IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of:

Itzhak Parnafes et al.

Serial No.: 09/586,531

Filed: May 31, 2000

For: METHOD AND APPARATUS PROVIDING AUTOMATIC RESV MESSAGE GENERATION FOR NON-RESV-CAPABLE NETWORK DEVICES

Confirmation No.: 6019

Group Art Unit No.: 2663

Examiner: Derrick W. Ferris

Mail Stop Petitions
Commissioner for Patents
P. O. Box 1450
Alexandria, VA 22313-1450

DECLARATION UNDER 37 C.F.R. 1.132

We, the inventors of the above-identified patent application, namely, Shai Mohaban, Itzhak Parnafes, Silvano Gai, and Dinesh Dutt, hereby declare:

1. We are inventors of the subject matter disclosed and claimed in the above-referenced patent application now pending in the United States Patent & Trademark Office ("Office"), as each of us contributed to the conception of at least one claim in the patent application. Any use of the word "our," "we," or "us" in this declaration refers collectively to the four named inventors of the patent application, namely Shai Mohaban, Itzhak Parnafes, Silvano Gai, and Dinesh Dutt.

2. We have been informed that the Office has rejected the claims of the application as allegedly anticipated by the Network Working Group Internet Draft entitled "RSVP Receiver Proxy" by Gai et al ("the *Gai* reference"), which is attached hereto as **Exhibit A**. We have read the *Gai* reference.

3. The purpose of this declaration is to establish that the contents of the *Gai* reference are derived from our own work.

4. Prior to October 1999, we conceived of the invention described and claimed in this application. Each of us contributed to the conception of at least one claim in the application.

Shai Mohaban and Itzhak Parnafes were co-founders of a company, Class Data, that was acquired by Cisco Systems, Inc. Silvano Gai and Dinesh Dutt were employed by Cisco Systems, Inc. at the time that Shai Mohaban and Itzhak Parnafes become employees of Cisco Systems, Inc. as a result of the acquisition. Prior to joining Cisco Systems, Inc., Shai Mohaban and Itzhak Parnafes were working on an approach for establishing a network resource reservation at a client. At the time Shai Mohaban and Itzhak Parnafes become employees of Cisco Systems, Inc., Silvano Gai and Dinesh Dutt were working on an approach for establishing a network resource reservation at a router. After Shai Mohaban and Itzhak Parnafes become employees of Cisco Systems, Inc., Shai Mohaban, Itzhak Parnafes Silvano Gai, and Dinesh Dutt collaborated together to develop the subject matter identified by the claims of the patent application.

5. As a general matter of practice, technology developers often discuss technical solutions with their peers. Such discussions are helpful as they often help assess the merit of, and gain support for, the technical solution. In this spirit, after we conceived of the invention described and claimed in this application, two of the inventors, namely Silvano Gai and Dinesh Dutt, discussed the invention as described and claimed in this application with Nitsan Elfassy, who was employed by Cisco Systems, Inc. at the time, and Yoram Bernet, who was employed by Microsoft Corporation at the time.

As a result of these discussions, Silvano Gai, Dinesh Dutt, Nitsam Elfassy, and Yoram Bernet authored the *Gai* reference. The purpose of the *Gai* reference was to present our

invention (i.e., the invention that was invented by Shai Mohaban, Itzhak Parnafes, Silvano Gai, and Dinesh Dutt) to the Internet community to solicit their comments.

Silvano Gai and Dinesh Dutt had experience in authoring Network Working Group Internet Drafts, as they had each previously authored at least one Network Working Group Internet Draft prior to the *Gai* reference. As it was not necessary to have all of the four inventors author the *Gai* reference, Shai Mohaban and Itzhak Parnafes were not involved in authoring the *Gai* reference; thus, Shai Mohaban and Itzhak Parnafes were not listed as an author of the *Gai* reference.

The contributions of Nitsan Elfassy and Yoram Bernet were limited to: (a) discussing applications of the invention, after the invention was conceived, with Silvano Gai and Dinesh Dutt, and (b) reviewing the *Gai* reference. Nitsan Elfassy and Yoram Bernet did not contribute to the conception of any claim listed in the patent application, and thus, do not qualify as inventors of the patent application. However, because of their work in discussing applications of the invention and reviewing the *Gai* reference, it was appropriate to identify Nitsan Elfassy and Yoram Bernet as co-authors of the *Gai* reference.

As a general matter of practice, it is desirable to list as authors of a Network Working Group Internet Draft individuals from more than one organization, since Network Working Group Internet Drafts authored by a single organization are given less weight by the technical community. In this spirit, Silvano Gai and Dinesh Dutt wished to include Yoram Bernet as a co-author of the *Gai* reference because Internet Drafts are typically presented by representatives of at least two organizations.

6. To the extent that the *Gai* reference discloses aspects of the invention, those aspects of the invention disclosed in the *Gai* reference either describe our own work, or are

derived from our own work. Therefore, the *Gai* reference describes our own work, or derives from our own work, as described and claimed in this application.

7. The undersigned being warned that willful false statements and the like are punishable by fine or imprisonment, or both, under 18 U.S.C. 1001, and that such willful false statements and the like may jeopardize the validity of the application or document or any patent resulting therefrom, declares that all statements herein made of his/her own knowledge are true, and all statements herein made on information and belief are believed to be true.

_____ Signed at _____, this ____ day of
SHAI MOHABAN _____, 2006.

_____ Signed at _____, this ____ day of
ITZHAK PARNAFES _____, 2006.

_____ Signed at _____, this ____ day of
SILVANO GAI _____, 2006.

_____ Signed at _____, this ____ day of
DINESH DUTT _____, 2006.

50325-0085 *Patent*

UNITED STATES PATENT APPLICATION

FOR

METHOD AND APPARATUS PROVIDING AUTOMATIC RESV MESSAGE GENERATION FOR NON-RESV-CAPABLE NETWORK DEVICES

INVENTORS:

ITZHAK PARNAFES
SHAI MOHABAN

PREPARED BY:

HICKMAN PALERMO TRUONG & BECKER LLP
1600 WILLOW STREET
SAN JOSE, CA 95125-5106
(408) 414-1080

"Express Mail" mailing label number EL624353128US

Date of Deposit 5/31/00

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.

CASEY MOORE
(Typed or printed name of person mailing paper or fee)

Casey Moore
(Signature of person mailing paper or fee)

METHOD AND APPARATUS PROVIDING AUTOMATIC RESV MESSAGE GENERATION FOR NON-RESV-CAPABLE NETWORK DEVICES

FIELD OF THE INVENTION

The present invention generally relates to computer networking. The invention relates more specifically to a method and apparatus for automatically facilitating reservations on network devices for a traffic flow directed from an originating node to a receiving node that is enabled to participate in such reservation.

BACKGROUND OF THE INVENTION

A computer network typically comprises a plurality of interconnected entities that transmit date frames ("sources," "senders") or entities that receive data frames ("sinks," "receivers," "destinations"). A common type of computer network is a local area network ("LAN") that generally comprises a privately owned network within a single building or campus. LANs employ a data communication protocol (LAN standard) such as Ethernet, FDDI, or Token Ring, that defines the functions performed by the data link and physical layers of a communications architecture (i.e., a protocol stack), such as the Open Systems Interconnection (OSI) Reference Model. In many instances, multiple LANs may be interconnected by point-to-point links, microwave transceivers, satellite hookups, etc., to form a wide area network ("WAN"), metropolitan area network ("MAN") or Intranet. These internetworks may be coupled through one or more gateways to the global, packet-switched internetwork known as the Internet.

Each network entity preferably includes network communication software, which may operate in accordance with Transport Control Protocol/Internet Protocol (TCP/IP) or some other suitable protocol. A protocol generally consists of a set of rules defining how entities interact with each other. In particular, TCP/IP defines a series of communication layers, including a transport layer and a network layer. At the transport layer, TCP/IP includes both the User Data Protocol (UDP), which is a connectionless transport protocol,

and TCP, which is a reliable, connection-oriented transport protocol. When a process at one network entity (source) wishes to communicate with another entity, it formulates one or more messages and passes them to the upper layer of the TCP/IP communication stack. These messages are passed down through each layer of the stack where they are encapsulated into packets and frames. Each layer also adds information in the form of a header to the messages. The frames are then transmitted over the network links as bits. At the destination entity, the bits are re-assembled and passed up the layers of the destination entity's communication stack. At each layer, the corresponding message headers are also stripped off, thereby recovering the original message which is handed to the receiving process.

One or more intermediate network devices are often used to couple LANs together and allow the corresponding entities to exchange information. For example, a bridge may be used to provide a "bridging" function between two or more LANs. Alternatively, a switch may be utilized to provide a "switching" function for transferring information, such as data frames or packets, among entities of a computer network. Typically, the switch is a computer having a plurality of ports (i.e., logical network interfaces ("LI" or "interfaces)) that couple the switch to several LANs and to other switches. The switching function includes receiving data frames at a source port and transferring them to at least one destination port for receipt by another entity. Switches may operate at various levels of the communication stack. For example, a switch may operate at Layer 2 which, in the OSI Reference Model, is called the data link layer, and includes the Logical Link Control (LLC) and Media Access Control (MAC) sub-layers.

Other intermediate devices, commonly known as routers, may operate at higher communication layers, such as Layer 3, which, in TCP/IP networks, corresponds to the Internet Protocol (IP) layer. IP data packets include a corresponding header which contains an IP source address and an IP destination address. Routers or Layer 3 switches may re-assemble or convert received data frames from one LAN standard (e.g., Ethernet)

to another (e.g., Token Ring). Thus, Layer 3 devices are often used to interconnect dissimilar sub-networks. Some Layer 3 intermediate network devices may also examine the transport layer headers of received messages to identify the corresponding TCP or UDP port numbers being utilized by the corresponding network entities. Many applications are assigned specific, fixed TCP and/or UDP port numbers in accordance with Request For Comments (RFC) 1700. For example, TCP/UDP port number 80 corresponds to the Hypertext Transport Protocol (HTTP), while port number 21 corresponds to File Transfer Protocol (FTP) service.

-- QUALITY OF SERVICE (QoS)

Networks that use the packet data communication technology as described above are now undergoing adaptation to carry voice traffic. In such "voice over IP" ("VoIP") networks, processes executing at network entities (e.g., Internet hosts) may generate hundreds or thousands of traffic flows that are transmitted across a network. Generally, a traffic flow is a set of messages (frames and/or packets) that typically correspond to a particular task, transaction, or operation (e.g., a print transaction), may be associated with an application, and may be characterized by values of various network and transport parameters such as source and destination IP addresses, source and destination TCP/UDP port numbers, and transport protocol.

Computer networks typically include numerous services and resources for use in moving traffic flows throughout the network. For example, different network links, such as Fast Ethernet, Asynchronous Transfer Mode (ATM) channels, network tunnels, satellite links, etc., offer unique speed and bandwidth capabilities. Particular intermediate devices also include specific resources or services, such as number of priority queues, filter settings, availability of different queue selection strategies, congestion control algorithms, etc. Each port/logical network interface (LI) of a network device can provide a different service or resource. For ease of explanation, the term "network device,"

unless expressly stated otherwise, herein refers to the device in its entirety or one or more ports/LIs.

To maximize the performance of a traffic flow across a network, a desired quality of service (QoS) can be requested. For a given traffic flow, a desired QoS can be designated for each of various aspects of the traffic flow treatment across the network (i.e., how various services and resources of the network interact with the traffic flow as it travels across the network). To deliver voice over IP traffic with acceptable quality, it is important to ensure that all network elements in a network path from sender to receiver are configured with QoS adequate to support VoIP.

-- RESERVATION OF NETWORK RESOURCES

To obtain desired qualities of service (QoS), known mechanisms can be used to reserve resources across a network along a path between a source and intended destination of an intended future traffic flow. For example, the Resource Reservation Protocol (RSVP) provides a mechanism by which such resource reservations can be made. RSVP is defined in Internet Request For Comment (RFC) 2205. The RSVP protocol is a mechanism to establish a network resource reservation along a path from a sender to a receiver (or multiple receivers, in the case of multicast). Generally, it requires both the sender and receiver to be actively engaged in establishing the reservation.

When a source is capable of utilizing RSVP (i.e., it is RSVP-enabled), the source can generate and send an RSVP Path message along an intended path to the destination (i.e., intended, i.e., anticipated receiver). The RSVP Path message can specify the various characteristics of the traffic flow that is to be sent. Since IP is a connectionless protocol, RSVP can be used to set up paths for a traffic flow and guarantee bandwidth on the paths.

If the destination device is RSVP-enabled, upon receipt of the RSVP Path message the destination device can ignore the RSVP Path message, raise an error condition, or apply for reservation of one or more resources along the intended path. This

can be done by generating and sending a RESV message through the intended path (including each device along that path) to the source. The RESV message can be recognized by RSVP-enabled devices along the intended path, which can either ignore the RESV message or reserve resources for the anticipated traffic flow.

Additional information regarding RSVP can be found in, e.g., U. Black, "Voice Over IP" (Upper Saddle River, NJ: Prentice-Hall PTR, 2000), at 210; "Cisco Internetworking Technologies Handbook" (Macmillan Technical Publishing, 1998). The RSVP mechanism can be used with multicast or unicast traffic flows, and the discussions herein are equally applicable to both with suitable modification.

Unfortunately, some destinations of traffic flows are not designed or are otherwise unable to participate in a reservation process by returning a RESV message as discussed above. For example, the destination may be non-RSVP-enabled, not trusted, or merely designed or utilized so as not to bear the burden of handling the issue. For ease of description, such destinations are hereinafter collectively referred to as non-RESV-capable.

Further, supporting RSVP signaling can require a network device to have a large and complex protocol stack, with heavy demands on the source with respect to the memory footprint, O/S capabilities, and CPU load. Omitting these elements from network devices, while still providing RSVP capability in some way, would reduce the cost of such devices and fulfill a market need.

Providing large volumes of such less expensive devices, especially consumer-oriented end devices, can be commercially attractive to device manufacturers. An example of such devices is the IP phone; IP phones are not expected to be RSVP enabled. However, deployment of end-to-end RSVP signaling may be crucial in ensuring that voice-over-IP connections of reasonable quality can be consistently established.

Still another problem of present approaches is that there are large numbers of deployed and installed devices that do not currently support RSVP and RESV. Re-

configuring or updating these devices to support such protocols would be expensive, time-consuming, and could require adding more memory or processing power to the existing devices, which is undesirable.

When applications or other destinations on such lower capability devices are the intended destination of a traffic flow, the RSVP mechanism of providing QoS for the traffic flow typically cannot be utilized. Unfortunately, to provide services, such as carrying voice over IP, with desirable quality, RSVP capability can be critical. Yet upgrading such devices to support RSVP can prohibitively increase the cost of the devices.

Based on the foregoing, there is a clear need in this field for a mechanism for generating and communicating a RESV message associated with a traffic flow that is intended to be sent to a non-RESV-capable destination.

In particular, there is a need to provide such a system and method with minimal cost and complexity and maximum efficiency.

There is a need to provide a way for non-RSVP-enabled devices to recognize RSVP messages and respond with network resources reservations, at minimum cost and without modifying such existing devices.

SUMMARY OF THE INVENTION

The foregoing objects and advantages, and other objects and advantages that will become apparent from the following description, are achieved by the present invention, which comprises, in one embodiment, a method of establishing a network resources reservation for an anticipated traffic flow along a path in a network between an anticipated source and an anticipated receiver of the traffic flow, wherein the anticipated receiver otherwise cannot facilitate establishing the network resources reservation. The method may comprise detecting an RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a proxy node located within the path, determining whether to establish the network resources reservation, generating an RESV message to reserve network resources for the anticipated traffic flow, and communicating the RESV message to the anticipated source of the anticipated traffic flow.

One feature involves determining one or more device and traffic parameter values associated with the anticipated traffic flow, and a related feature is that the step of generating the RESV message comprises the step of generating the RESV message based on at least one of the device and traffic parameter values.

In another feature, the method further comprises receiving predefined policy information and generating the RESV message based on the predefined policy information. In yet another feature, the step of determining whether to initiate an RSVP reservation process includes the steps of determining one or more network parameter values associated with the anticipated traffic flow, determining one or more transport parameter values associated with the anticipated traffic flow, determining next and previous hop parameter values associated with the anticipated traffic flow, and correlating at least one of the ascertained network parameter, transport parameter, next hop parameter, and previous hop parameter values with information defining a relationship between them and whether a RESV message is desired. In yet another feature, the step of detecting an RSVP Path message comprises the step of detecting an RSVP Path message

associated with the anticipated receiver of the anticipated traffic flow at a proxy node that is logically positioned adjacent to the path.

In another aspect, a proxy device is located on the path from the sender to the receiver, in a network in which the sender uses RSVP signaling. The proxy device receives the RSVP Path message, inspects the message, and determines whether it should serve as an auto-reserve agent or RESV proxy on behalf of the receiver. The decision can be based on parameters found in the Path message, such as session parameters (e.g., source and destination IP addresses and port numbers, and protocol), next hop and previous hop values, and policy information that is carried in the message. If the proxy device decides to serve as an auto-reserve agent, then the proxy device returns an RESV message to the sender on behalf of the receiver. The proxy device may forward the Path message to the receiver, or not forward it and thereby serve as a terminator of it. In the cause of multicast signaling, the proxy device can send RESV messages for only some of all the intended receivers, and may forward the Path message to a subset of receivers.

Other features and aspects will become apparent from the following detailed description.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings in which like reference numerals refer to similar elements and in which:

FIG. 1A is a block diagram of a first network including a RESV proxy, in accordance with an embodiment of the present invention.

FIG. 1B is a block diagram of a second network including a RESV proxy, in accordance with an embodiment of the present invention.

FIG. 2 is a schematic representation of constituents of a RESV proxy, in accordance with an embodiment of the present invention.

FIG. 3 is a flow diagram of a method for generating and communicating a RESV message, for a traffic flow targeted to a non-RESV-capable destination, in accordance with an embodiment of the present invention.

FIG. 4 is a flow diagram of a method that can be included as part of the method of FIG. 3, in accordance with an embodiment of the present invention.

FIG. 5 is a flow diagram of operations that can be included within a portion of method in FIG. 3.

FIG. 6 is a block diagram of a computer system with which an embodiment of the present invention may be carried out.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

In one embodiment, a method and apparatus is described for generating and communicating a RESV message in response to an RSVP Path message associated with a traffic flow intended to be sent to a non-RESV-capable destination. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

-- STRUCTURAL OVERVIEW OF RESV PROXY SYSTEM

FIG. 1A is a block diagram of an exemplary computer network that is configured to facilitate the reservation of network resources, according to RSVP, for a traffic flow that is intended to be received by a destination that is not capable of facilitating such a reservation itself.

Network 100 includes a sender node (i.e., source of an RSVP Path message) 102, a RESV proxy 104, one or more receiver nodes (i.e., destinations) 106, one or more network devices 108, a policy server 110, and the remainder of the network, or sub-network 112 which may include one or more additional network devices. Each of these components is connected to each other, directly or indirectly, physically or otherwise.

In particular, RESV proxy 104 is connected in the network path in a position where it will receive information passed from the sender node 102 and intended to be received by the receiver node 106 (either physically in the path as shown here, or able to sniff traffic, as in FIG. 1B described below). As an example, in FIG. 1A the RESV proxy 104 is shown with no network devices between it and the receiver node 106; however, one or more network devices can be connected between them. Similarly, none or any number of network devices can be connected between the RESV proxy 104 and the sender node 102.

The sender node 102 is any suitable network device that includes a source that can generate an RSVP Path message associated with an anticipated traffic flow, to initiate an associated resource reservation process (RSVP-enabled). Such a source can be, for example, a client or server computer that executes the Windows 2000 operating system, a router, VoIP Gateway, etc. The source can be configured to also generate the anticipated traffic flow, and/or can be an RSVP proxy.

One or more network devices 108 and the remainder of the network, or sub-network 112 can be located along a send traffic flow path (indicated by arrows Ps) between the sender node 102 and the one or more receiver nodes 106. One or more of the network devices 108 and the devices of sub-network 112 can be configured to reserve resources in the network or in themselves in accordance with RSVP.

The one or more receiver nodes 106 can be any suitable network devices or applications that can be the intended destination device of a traffic flow originated by the sender node 102. Typically, a receiver node is a server that contains network resources needed by a client in the position of the sending node. If the traffic flow is unicast, only one receiver node 106 will be at the end of the traffic flow path P, while more than one receiver node 106 may be included if the traffic flow is multicast. At least one of the receiver nodes 106 can be incapable of responding to the receipt of an RSVP Path message by generating and communicating a RESV message (i.e., non-RESV-capable).

RESV proxy 104 is located in the network so that it can detect new RSVP Path messages sent by the sender node 102. In FIG. 1A, this location is on the traffic flow path Ps (i.e., as an intermediate gateway in the message's route) and the message passes to the RESV proxy. For example, the RESV proxy 104 can be located at the first switch or router to which the receiver node 106 is connected.

Such a device can also be configured to detect various parameter values in the RSVP Path message that are associated with the network, and transport parameter values associated with the anticipated traffic flow. Such parameter values include source and

destination IP addresses and port numbers, and the transport protocol. The RSVP protocol identifies SESSION and SENDER_TEMPLATE objects that carry these values, as stated in RFC 2205.

RESV proxy 104 is also configured to determine whether it is desirable to generate and communicate a RESV message, associated with the detected RSVP Path message, back to the sender node 102 along the RESV message flow path (indicated by the arrows Pr). Such configuration can include predetermined relationships or logic based on characteristics of the RSVP Path message, such as information regarding the intended receiver node of the traffic flow associated with the RSVP Path message. Information regarding the intended receiver node is specified in the SCOPE object of an RSVP message. Such characteristics or parameters can then be used to determine whether to generate and communicate RESV signaling for the flow.

For example, the RESV proxy can be configured, for example by a network administrator, to generate and communicate a RESV message when characteristics of or objects in a detected RSVP Path messages are associated with anticipated traffic flows directed to a particular receiver node, or a type of receiver node.

In some embodiments, the RESV proxy might not be configured to determine whether a RESV message is needed. For example, such capability may not be needed if the RESV proxy encounters RSVP Path flows that are intended to be received by only receiver nodes 106 that are non-RESV-capable. In such cases, it is desirable for the RESV proxy to always generate a RESV message. Using such a RESV proxy would benefit from appropriate placement within the network, and/or appropriate configuration, such that the RESV proxy encounters RSVP Path messages intended to be received by only receiver nodes that are non-RESV-capable.

As is shown in FIG. 1A, a policy server 110 is located on the network, but is not necessarily located along the send traffic flow path Ps. The policy server 110 can be any suitable device that can include one or more quality of service policies associated with

one or more attributes of traffic flows and/or other network attributes. Thus, policy server 110 can include predetermined relationships or logic that can define an appropriate policy or policies in accordance with the one or more attributes of a given traffic flow. An example of policy server 110 is Cisco Quality of Service Policy Manager.

Policy server 110 can be configured to return an appropriate policy for generating and communicating the RESV message, including defining characteristics of the RESV message, based on the characteristics of the RSVP Path message associated with the traffic flow and/or the characteristics of the intended receiver node(s). For example, in the case of a flow representing a voice call, a policy can be predefined indicating that all routers must be configured to use immediate packet forwarding or guaranteed bandwidth for that flow.

Policy server 110 can be dedicated to the RESV proxy operation, or can also provide policies for other aspects of the network 100 operation. While the policy server 110 is depicted in FIG. 1A as a device separate from the RESV proxy 104, the policy server 110 or its equivalent alternatively or additionally can reside on the RESV proxy device 104.

FIG. 1B depicts a block diagram of a network 200. Like numbered elements of FIG. 1B as compared to FIG. 1A are substantially the same as those described above with reference to FIG. 1.

In FIG. 1B, however, RESV proxy 204 is connected to a shared medium along the path Ps, rather than being located directly along the send traffic flow path Ps. For example, the RESV proxy 204 can be located at an external host on the same shared media (e.g., Ethernet) as the sender node, or such a host connected to a spanning port of the switch to which the sender node is connected. Alternatively, the RESV proxy 204 can be located at an ingress point into a different network (e.g., into an ISP, or to a WAN connection).

In this configuration, RESV proxy 204 can operate as a packet sniffer to detect new RSVP Path messages that pass along the send traffic flow path Ps, as well as characteristics of the RSVP Path message associated with the anticipated traffic flow. The RESV proxy 204 is also configured to determine whether it is desirable to generate and communicate a RESV message, associated with the detected RSVP Path message and anticipated traffic flow, along the RESV traffic flow path Pr. Such configuration can include predetermined relationships and/or logic correlating characteristics of the RSVP Path message or characteristics of the traffic flow derived from the RSVP Path message characteristics.

For example, RESV proxy 204 may be configured on a host of a subnet of network 200. In this position, RESV proxy 204 listens to all RSVP messages entering the subnet and issues RESV messages on behalf of devices in the subnet, where the subnet hosts are not capable of recognizing or responding to RSVP messages.

An RESV proxy as described herein provides the benefits of RSVP resource reservation to end devices that are not themselves configured to facilitate such a process by responding to a RSVP Path message with a RESV message. In addition, using one such device in a location separate from the end device, these benefits are provided inexpensively and without increasing the complexity or cost of the end device. Furthermore, a single such device can be included in the paths of traffic flows from more than one sender device, thus magnifying the simplification and cost savings over what would be required to provide the same capabilities directly to the various end devices themselves.

FIG. 2 is a schematic representation of elements that comprise RESV proxy 104 of FIG. 1A, according to one embodiment. As shown in FIG. 2, RESV proxy 104 includes an RSVP Path message detector 250, an RSVP Path message analyzer 252, and a RESV message generator 254.

RSVP Path message detector 250 is configured to detect signals that form an RSVP Path message. The RSVP Path message analyzer 252 is configured to identify and recognize the signals of an RSVP Path message detected by the RSVP Path message detector 250. These identified and recognized signals are then used by the RESV message generator 254, which is configured to generate signals that form a RESV message associated with, and based on, the RSVP Path message signals.

In one embodiment, RESV proxy 204 of FIG. 1B also comprises the elements shown in FIG. 2, appropriately modified to carry out detection of the RSVP Path messages and returning of the generated RESV message.

The elements of FIG. 2 may comprise one or more programmatic objects, methods, processes, subroutines, or other software elements that individually or collectively carry out the operational functions disclosed herein. In one embodiment, the elements of FIG. 2 are implemented in a router, gateway, or other network device. In another embodiment, the elements of FIG. 2 are executed by a server in the position of a receiver node. In still another embodiment, the elements of FIG. 2 are executed by a separate server or computer.

-- FUNCTIONAL OVERVIEW

FIG. 3 is a block diagram of an exemplary method 300 for responding to an RSVP Path message associated with an anticipated traffic flow with at least one anticipated receiver, where the anticipated receiver is not itself capable of appropriately responding to the RSVP Path message.

An RSVP path message is detected in operation 302. In one embodiment, operation 302 involves detecting the message directly, or indirectly, e.g., with a packet sniffer. As described above, the RSVP Path message can include one or more traffic flow parameter values associated with the anticipated future traffic flow. The traffic flow parameters can include, for example, source and destination IP addresses and port numbers, protocol, next and previous hop parameters, policy elements, etc.

In operation 304, a decision is made whether a path reservation should be initiated, i.e., whether a RESV message should be generated and communicated. For example, if the RSVP Path message is associated with more than one intended receiver (e.g., the anticipated traffic flow is anticipated to be a multicast), operation 304 includes a decision for each such intended receiver.

The decision whether to establish a path reservation in operation 304 can be based upon one or more traffic flow parameter values included in the RSVP Path message. For example, the Path message may include a FLOWSPEC object or a FILTER_SPEC object that carry values indicating quality of service parameters that the receiving device can implement. In such a case, the decision whether to establish a path reservation may depend on whether the values of such objects specify a traffic flow that needs resources to be reserved.

Alternatively, or additionally, the decision of operation 304 can be based upon a previous configuration of the RESV proxy, or based on one or more predefined policies specifying logic, relationships, or rules for determining what type of resource reservation is appropriate. Such policies can reside, for example, in a policy server within the network or within the RESV proxy itself.

Different decisions of whether to establish path reservations for each of more than one intended receiver can be based on different traffic flow parameter information associated with each of such receivers, or based on differences between the receivers. For example, information about such differences is based on previous RESV proxy configuration and/or received policies. In some instances, although the anticipated receiver of the RSVP Path message is capable of generating and communicating a RESV message, it can nevertheless be decided in operation 304 that a RESV message should be generated.

If operation 304 determines that a path reservation should not be established (for example, if the intended receiver is RSVP-enabled and can, therefore, respond to the

RSVP Path message), then the RSVP Path message is forwarded, as shown by operation 306. The RSVP Path message can be forwarded to the anticipated receiver (i.e., destination) of the traffic flow associated with the RSVP Path message detected in operation 302. If the RSVP Path message is associated with more than one intended receiver (e.g., the traffic flow will be a multicast), then the RSVP Path message can be forwarded in operation 306 to one or more of the multiple anticipated receivers for which it is determined that a path reservation should not be established by the RESV proxy.

In alternative embodiments, method 300 does not include operation 304 or operation 306. Instead, control is passed from operation 302 to operation 308, in which an RESV message is generated, as described below.

If operation 304 determines that path reservation should be generated, and in alternative embodiments where operation 304 is omitted, a RESV message is generated in operation 308. The content of the RESV message can be based, at least in part, on attributes of the detected RSVP Path message. For example, each Path message normally contains a SENDER_TEMPLATE object defining the format of the data packets and a SENDER_TSPEC object specifying the traffic characteristics of the flow. The resources needed for a reservation, and the content of an RESV message, can be based on the values of these objects.

Further, message attributes or object values of an RSVP Path message can be correlated with predefined relationships or logic, such as policies, that relate such object values or attributes to RESV message components. The predefined relationships can be defined prior to the detection of the RSVP Path message of method 300, for example, during configuration of the RESV proxy.

As a result, an RESV message is generated that includes information identifying the requested resource reservation in network devices along the intended path between the source and destination. If resources matching the parameters of the RSVP Path message are not available in the network path from sender to receiver, then the requested

resource reservation can differ from that suggested according to the RSVP path message detected in operation 302.

In operation 310, the resulting RESV message is communicated. In one embodiment, operation 310 involves sending the RESV message along the intended path, as defined based on the RSVP Path message detected in operation 302. Operation 310 also includes installing a RESV state on each of the network devices along the intended path, and returning the RESV message to the source, on behalf of the anticipated receiver or receivers.

FIG. 4 is a block diagram of a method that can be included in the method 300 of Figure 3, for example, between operations 304 and operation 308.

Method 320 of FIG. 4 includes determining, in operation 322, whether to forward the RSVP Path message to one or more intended receivers, even though operation 304 has determined to generate and communicate a RESV message. Operation 322 may be used, for example, when an RESV proxy in the position of receiver wishes to issue a counter-message that defines a different set of traffic characteristics for the flow initiated by the sender. Operation 322 also may be used when the RESV proxy wishes to explicitly announce its presence in the network to an RSVP-enabled device that is in the position of sender.

If operation 322 determines not to forward the RSVP Path message, then method 320 returns control or ends. For example, in the case where method 320 is placed between operations 304 and 308 of method 300 in FIG. 3, method 300 will continue with operation 308.

If operation 322 determines to forward the RSVP Path message, then the RSVP Path message is forwarded in operation 324. As with operation 304 and operation 306 of method 300 described above, the determination of operation 322 can be made for each of one or more of the anticipated receivers. Thus, even if a RESV message is generated and

communicated in operation 308 and operation 310 for an anticipated receiver, the RSVP Path message may still be forwarded to that receiver.

FIG. 5 is a block diagram of operations that can be included in operation in 304 of method 300.

In operation 340, the anticipated traffic flow network and transport parameter values are determined from the RSVP Path message. As described above, the network and transport parameters can include, for example, the source and destination IP addresses and port numbers, and protocols. The network and transport parameters can further include the bandwidth, packet size, packet rate, and average rate.

In operation 342, the next and previous hop parameter values associated with the anticipated traffic flow are determined from the RSVP Path message. As defined in the RSVP protocol, a Path message may include an RSVP_HOP object that carries the IP address of the RSVP-capable node that sent the message and a logical outgoing interface handle. As stated in RFC 2205, an RSVP_HOP object may be a PHOP (previous hop) object for downstream messages or an NHOP (next hop) object for upstream objects.

Also, one or more policies associated with the RSVP Path message can be determined in operation 344. The RSVP protocol defines a POLICY_DATA object that may carry policy information. Thus, policy information can be obtained by extracting any POLICY_DATA object that forms part of the RSVP Path message. Such policies can define traffic flow treatment, as discussed above.

The various RSVP Path message attributes, e.g., traffic flow parameter values, are correlated in operation 346 with predefined relationships and/or logic relating such attributes with whether or not generation and communication of a RESV message is appropriate. For example, such relationships and/or logic can be in the form of policies such as "If FLOWSPEC = 3, then generate a RESV message" or "If FLOWSPEC = 1, then do not generate a RESV message." As defined in RFC 2205, FLOWSPEC=3 indicates a request for guaranteed quality of service, whereas FLOWSPEC=1 is a flow

specification requiring controlled delay. The predefined relationships or logic can be defined prior to the detection of the RSVP Path message of method 300, for example, during configuration of the RESV proxy. The correlation of operation 346 can be performed in conjunction with a policy server or other external policy decision point. Based on the information provided by the RESV proxy and policies resident on the policy server, the policy server can return the relationships or logic to the RESV proxy, or can make and return the determination of whether an RSVP Path message is needed.

If an embodiment of method 300 is performed by an RESV proxy, then the RESV message thereafter is communicated only through the sub-path between the proxy and the sender, or between a sniffing location along the sender-anticipated receiver path and the sender. In such a case, the ensuing path reservations will be made only in the network devices of that sub-path, and not in the network devices between the RESV proxy and the anticipated receiver. Therefore, it is desirable to locate the RESV proxy as close as possible to the anticipated receiver device. In alternative embodiments, with appropriate modification, the RESV proxy can generate a pseudo-RESV message and communicate the pseudo-RESV message along the remainder of the anticipated path between the RESV proxy and the anticipated receiver, thereby making a path reservation along more of the anticipated path.

The present invention provides a simple, efficient, low-cost mechanism for taking advantage of the benefits of RSVP in passing traffic flows through a network, even when the anticipated receivers of such traffic flows are not designed, or are otherwise unable, to facilitate a reservation process by responding to a RSVP Path message by generating and communicating a RESV message. Thus, RSVP resource scheduling can be used with traffic flows intended for receivers that are non-RSVP-enabled, not trusted, or spared the burden of such tasks, by design or usage. Furthermore, some devices can be less expensive when non-RSVP-enabled. The present invention, therefore, allows the use of such less expensive devices while still providing desired QoS. In addition, when a RESV

proxy device is included in the paths of traffic flows from more than one sender device, the complexity and cost savings over what would be required to provide the same capabilities directly to the various end devices themselves, are thereby magnified.

-- HARDWARE OVERVIEW

FIG. 6 is a block diagram that illustrates a computer system 600 upon which an embodiment of the invention may be implemented. Computer system 600 includes a bus 602 or other communication mechanism for communicating information, and a processor 604 coupled with bus 602 for processing information. Computer system 600 also includes a main memory 606, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 602 for storing information and instructions to be executed by processor 604. Main memory 606 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 604. Computer system 600 further includes a read only memory (ROM) 608 or other static storage device coupled to bus 602 for storing static information and instructions for processor 604. A storage device 610, such as a magnetic disk or optical disk, is provided and coupled to bus 602 for storing information and instructions.

Computer system 600 may be coupled via bus 602 to a display 612, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 614, including alphanumeric and other keys, is coupled to bus 602 for communicating information and command selections to processor 604. Another type of user input device is cursor control 616, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 604 and for controlling cursor movement on display 612. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

The invention is related to the use of computer system 600 for providing RSVP process facilitation for anticipated traffic flow receivers that do not otherwise do so alone.

According to one embodiment of the invention, providing RSVP process facilitation for anticipated receivers that do not otherwise do so alone is facilitated by computer system 600 in response to processor 604 executing one or more sequences of one or more instructions contained in main memory 606. Such instructions may be read into main memory 606 from another computer-readable medium, such as storage device 610. Execution of the sequences of instructions contained in main memory 606 causes processor 604 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

The term "computer-readable medium" as used herein refers to any medium that participates in providing instructions to processor 604 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 610. Volatile media includes dynamic memory, such as main memory 606. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 602. Transmission media can also take the form of acoustic or light waves, such as those generated during radio wave and infrared data communications.

Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punch cards, paper tape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 604 for execution. For example,

the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 600 can receive the data on the telephone line and use an infrared transmitter to convert the data to an infrared signal. An infrared detector can receive the data carried in the infrared signal and appropriate circuitry can place the data on bus 602. Bus 602 carries the data to main memory 606, from which processor 604 retrieves and executes the instructions. The instructions received by main memory 606 may optionally be stored on storage device 610 either before or after execution by processor 604.

Computer system 600 also includes a communication interface 618 coupled to bus 602. Communication interface 618 provides a two-way data communication coupling to a network link 620 that is connected to a local network 622. For example, communication interface 618 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 618 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 618 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

Network link 620 typically provides data communication through one or more networks to other data devices. For example, network link 620 may provide a connection through local network 622 to a host computer 624 or to data equipment operated by an Internet Service Provider (ISP) 626. ISP 626 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 628. Local network 622 and Internet 628 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 620 and through communication

interface 618, which carry the digital data to and from computer system 600, are exemplary forms of carrier waves transporting the information.

Computer system 600 can send messages and receive data, including program code, through the network(s), network link 620 and communication interface 618. In the Internet example, a server 630 might transmit a requested code for an application program through Internet 628, ISP 626, local network 622 and communication interface 618. In accordance with the invention, one such downloaded application provides for policy-based management of quality of service treatments of network data traffic flows by integrating policies with application programs as described herein.

The received code may be executed by processor 604 as it is received, and/or stored in storage device 610, or other non-volatile storage for later execution. In this manner, computer system 600 may obtain application code in the form of a carrier wave.

-- SCOPE

A method and apparatus providing RSVP resource reservation for anticipated traffic flows that are anticipated to be received by devices that do not facilitate the RSVP process have been described. The invention has been described herein in terms of several specific embodiments. Other embodiments of the invention, including alternatives, modifications, permutations and equivalents of the embodiments described herein, will be apparent to those skilled in the art from consideration of the specification, study of the drawings, and practice of the invention. For example, while only one RESV proxy has been shown, more than one RESV proxy can be included in a network. Also, the functionality of the RESV proxy can be included in the sender device, rather than a separate physical entity on the network. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. The embodiments and specific features described above in the specification and shown in the drawings should be considered exemplary rather than restrictive, with the invention being defined by the appended claims, which therefore include all such alternatives, modifications,

permutations and equivalents as fall within the true spirit and scope of the present invention.

CLAIMS

What is claimed is:

1. A method of establishing a network resources reservation for an anticipated traffic flow along a path in a network between an anticipated source and an anticipated receiver of the traffic flow, wherein the anticipated receiver otherwise cannot facilitate establishing the network resources reservation, the method comprising the steps of:

detecting an RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a proxy node located within the path;

determining whether to establish the network resources reservation;

generating an RESV message to reserve network resources for the anticipated traffic flow; and

communicating the RESV message to the anticipated source of the anticipated traffic flow.

2. A method as recited in claim 1, further comprising the step of determining one or more device and traffic parameter values associated with the anticipated traffic flow, and wherein the step of generating the RESV message comprises the step of generating the RESV message based on at least one of the device and traffic parameter values.

3. A method as recited in claim 1, further comprising the steps of:

receiving predefined policy information;

generating the RESV message based on the predefined policy information.

4. A method as recited in claim 1, wherein the step of determining whether to initiate an RSVP reservation process includes the steps of:

determining one or more network parameter values associated with the anticipated traffic flow;

determining one or more transport parameter values associated with the anticipated traffic flow;

determining next and previous hop parameter values associated with the anticipated traffic flow; and

correlating at least one of the ascertained network parameter, transport parameter, next hop parameter, and previous hop parameter values with information defining a relationship between them and whether a RESV message is desired.

5. A method as recited in claim 4, wherein determining the network parameter values and ascertaining the transport parameter values includes the steps of determining at least one of the source and receiver IP addresses, source and receiver port numbers, and transport protocol based on values carried in objects in the RSVP Path message.

6. A method as recited in claim 4, wherein determining the anticipated traffic flow characteristics includes determining at least one of the rate and size of packets associated with the anticipated traffic flow.

7. A method as recited in claim 4, further comprising the steps of extracting one or more additional anticipated traffic flow attributes from the RSVP Path message.

8. A method as recited in claim 7, wherein the anticipated receiver is an IP phone, and further comprising the step of determining at least one quality of service parameter associated with the anticipated traffic flow.

9. A method as recited in claim 1, further comprising the steps of:

communicating the RESV message along at least a subset of an anticipated path defined, at least in part, by the RSVP Path message;

receiving the RSVP Path message at one or more devices along the anticipated path.

10. A method as recited in claim 1, wherein the step of detecting an RSVP Path message comprises the step of detecting an RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a proxy node that is logically positioned adjacent to the path.

11. A computer readable medium comprising one or more sequences of instructions for facilitating an RSVP reservation process, for an anticipated traffic flow anticipated to be received by an anticipated receiver that cannot facilitate an RSVP reservation process for the anticipated traffic flow, wherein when the instructions are executed by one or more processors, the instructions cause the one or more processors to carry out the steps of:

detecting an RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a proxy node located within the path;

determining whether to establish the network resources reservation;

generating an RESV message to reserve network resources for the anticipated traffic flow; and

communicating the RESV message to the anticipated source of the anticipated traffic flow.

12. A computer-readable medium as recited in claim 11, further comprising the step of determining one or more device and traffic parameter values associated with the anticipated traffic flow, and wherein the step of generating the RESV message comprises the step of generating the RESV message based on at least one of the device and traffic parameter values.

13. A computer-readable medium as recited in claim 11, further comprising the steps of:

receiving predefined policy information;

generating the RESV message based on the predefined policy information.

14. A computer-readable medium as recited in claim 11, wherein the step of determining whether to initiate an RSVP reservation process includes the steps of:

determining one or more network parameter values associated with the anticipated traffic flow;

determining one or more transport parameter values associated with the anticipated traffic flow;

determining next and previous hop parameter values associated with the anticipated traffic flow; and

correlating at least one of the ascertained network parameter, transport parameter, next hop parameter, and previous hop parameter values with information defining a relationship between them and whether a RESV message is desired.

15. A computer-readable medium as recited in claim 14, wherein determining the network parameter values and ascertaining the transport parameter values includes the steps of determining at least one of the source and receiver IP addresses, source and receiver port numbers, and transport protocol based on values carried in objects in the RSVP Path message.

16. A computer-readable medium as recited in claim 14, wherein determining the anticipated traffic flow characteristics includes determining at least one of the rate and size of packets associated with the anticipated traffic flow.

17. A computer-readable medium as recited in claim 14, further comprising the steps of extracting one or more additional anticipated traffic flow attributes from the RSVP Path message.

18. A computer-readable medium as recited in claim 17, wherein the anticipated receiver is an IP phone, and further comprising the step of determining at least one quality of service parameter associated with the anticipated traffic flow.

19. A computer-readable medium as recited in claim 11, further comprising the steps of:

communicating the RESV message along at least a subset of an anticipated path defined, at least in part, by the RSVP Path message;

receiving the RSVP Path message at one or more devices along the anticipated path.

20. A computer-readable medium as recited in claim 11, wherein the step of detecting an RSVP Path message comprises the step of detecting an RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a proxy node that is logically positioned adjacent to the path.

21. A system for establishing a network resources reservation for an anticipated traffic flow along a path in a network between an anticipated source and an anticipated receiver of the traffic flow, wherein the anticipated receiver otherwise cannot facilitate establishing the network resources reservation, the system comprising:

means for detecting an RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a proxy node located within the path;

means for determining whether to establish the network resources reservation;

means for generating an RESV message to reserve network resources for the anticipated traffic flow; and

means for communicating the RESV message to the anticipated source of the anticipated traffic flow.

22. A network device that can establish a network resources reservation for an anticipated traffic flow along a path in a network between an anticipated source and an anticipated receiver of the traffic flow, wherein the anticipated receiver otherwise cannot facilitate establishing the network resources reservation, the network device comprising:

a network interface;

a processor coupled to the network interface and receiving network messages from the network through the network interface;

a computer-readable medium comprising one or more stored sequences which, when executed by the processor, cause the processor to carry out the steps of:

detecting an RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a proxy node located within the path;

determining whether to establish the network resources reservation;

generating an RESV message to reserve network resources for the anticipated traffic flow; and

communicating the RESV message to the anticipated source of the anticipated traffic flow.

ABSTRACT OF THE DISCLOSURE

A method and apparatus for providing network resource reservation capability for receiver nodes that either cannot or do not facilitate RSVP processes is provided. A RESV proxy is connected between an anticipated sender node and an anticipated receiver node in a computer network, e.g., a voice over IP network. The RESV proxy can detect an RSVP Path message and determine whether a RESV message is appropriate for the anticipated traffic flow. If it is so determined, the RESV proxy generates and communicates a RESV message through the network back to the anticipated sender node. The appropriateness of a RESV message can be based upon predetermined relationships or logic involving one or more of network parameters, transport parameters, and characteristics of the anticipated traffic flow, and other traffic flow attributes determined from analyzing the RSVP Path message. The RESV proxy can also interface with a policy server on the network to facilitate generation of the RESV message according to one or more of the anticipated traffic flow attributes. Alternatively, the RESV proxy can include the policy server functionality itself.

FIG. 1A

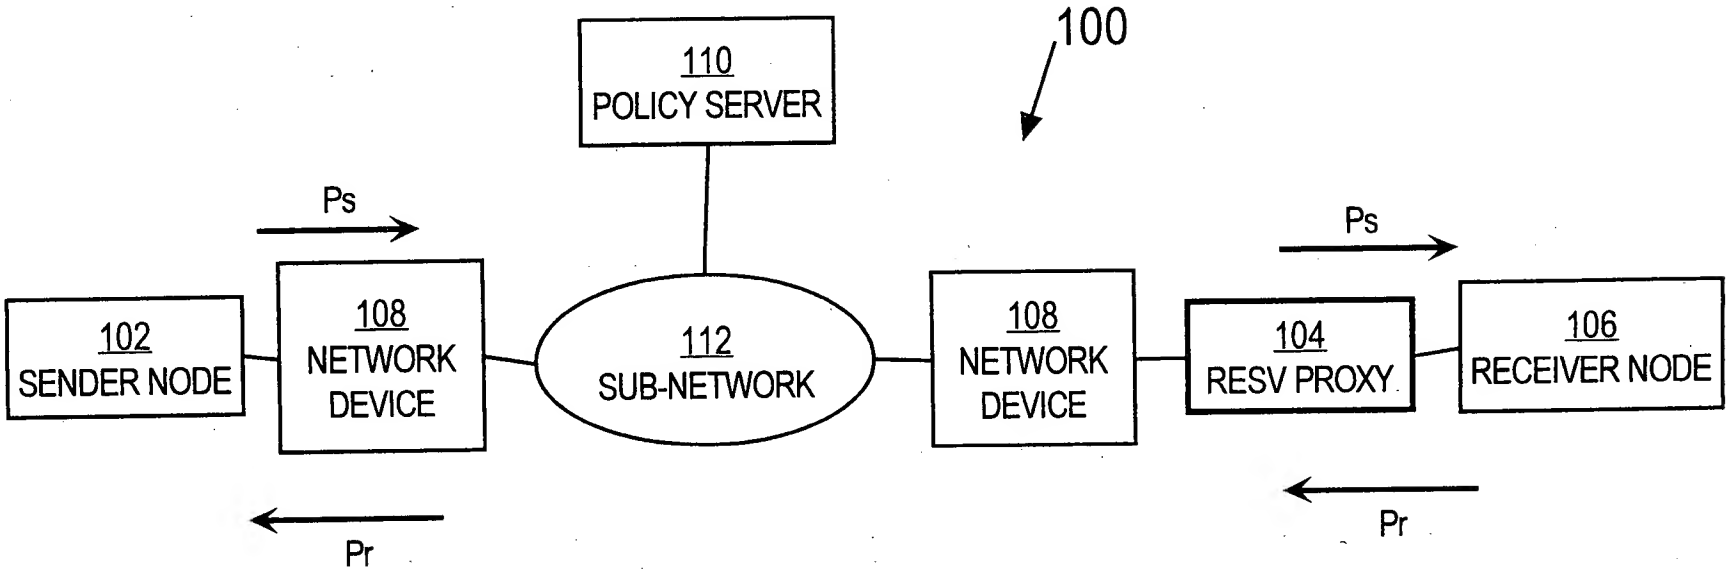
100
110
POLICY SERVER
Ps
Ps
102
SENDER NODE
108
NETWORK
DEVICE
112
SUB-NETWORK
108
NETWORK
DEVICE
104
RESV PROXY
106
RECEIVER NODE
Pr
Pr

FIG. 1B

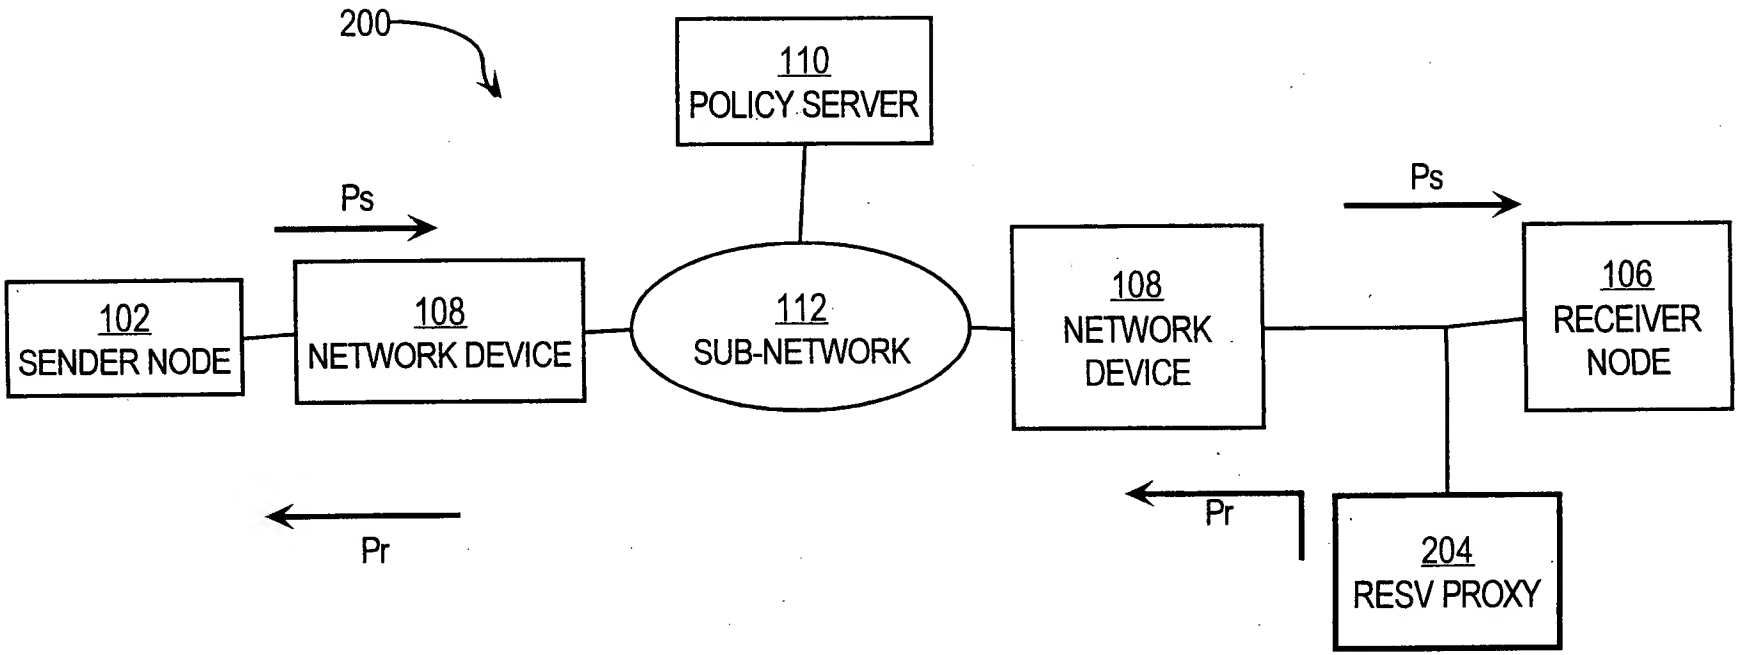
200
110
POLICY SERVER
Ps
Ps
102
SENDER NODE
108
NETWORK DEVICE
112
SUB-NETWORK
108
NETWORK
DEVICE
106
RECEIVER
NODE
Pr
Pr
204
RESV PROXY

FIG. 2

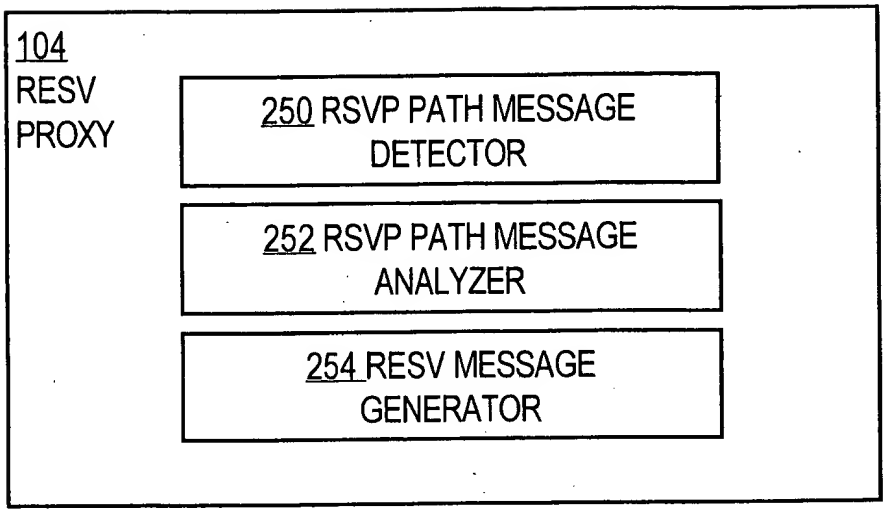
104
RESV
PROXY
250 RSVP PATH MESSAGE DETECTOR
252 RSVP PATH MESSAGE ANALYZER
254 RESV MESSAGE GENERATOR

FIG. 3

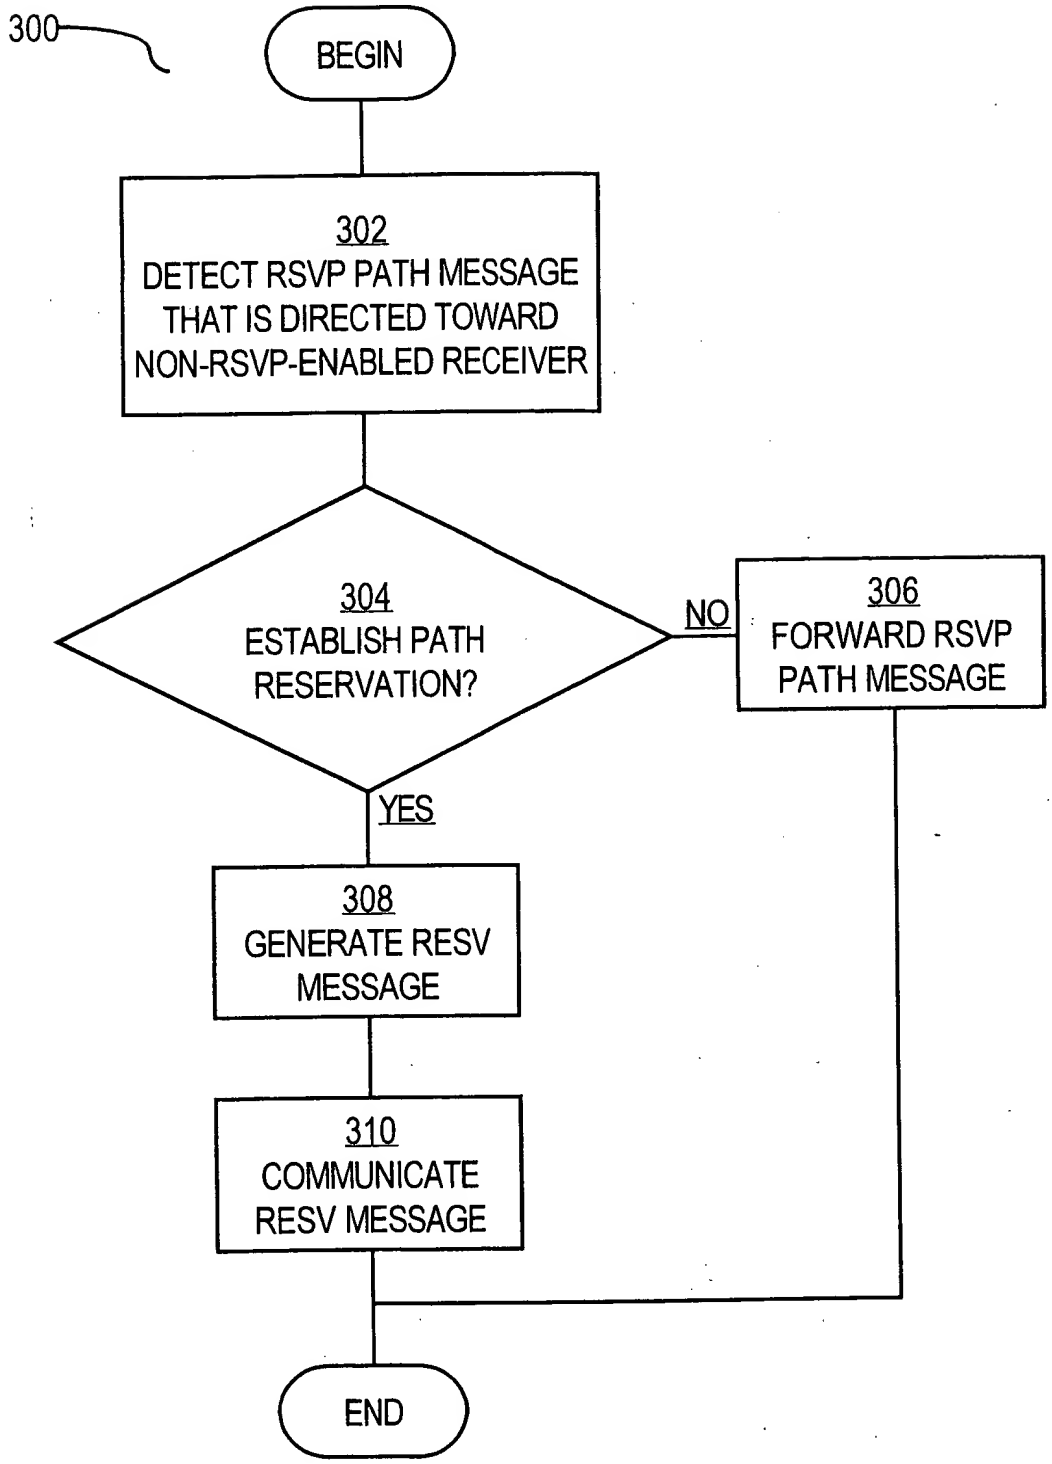
300
BEGIN
302
DETECT RSVP PATH MESSAGE THAT IS DIRECTED TOWARD NON-RSVP-ENABLED RECEIVER
304
ESTABLISH PATH RESERVATION?
NO
306
FORWARD RSVP PATH MESSAGE
YES
308
GENERATE RESV MESSAGE
310
COMMUNICATE RESV MESSAGE
END

FIG. 4

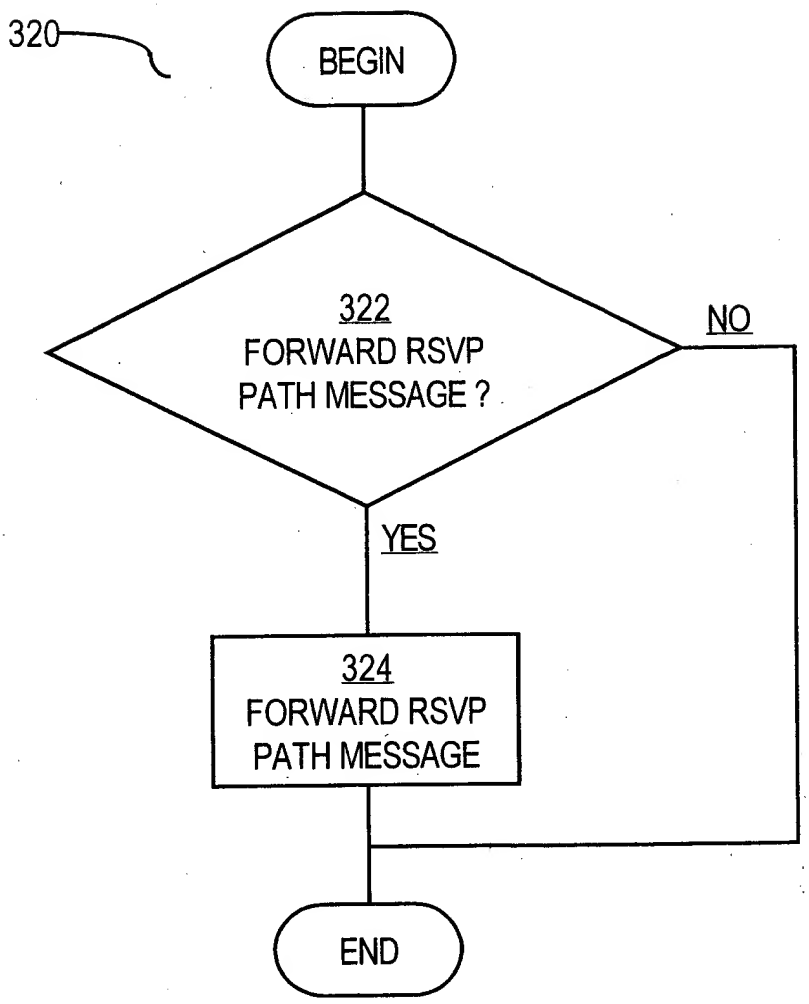
320
BEGIN
322
FORWARD RSVP
PATH MESSAGE ?
NO
YES
324
FORWARD RSVP
PATH MESSAGE
END

FIG. 5

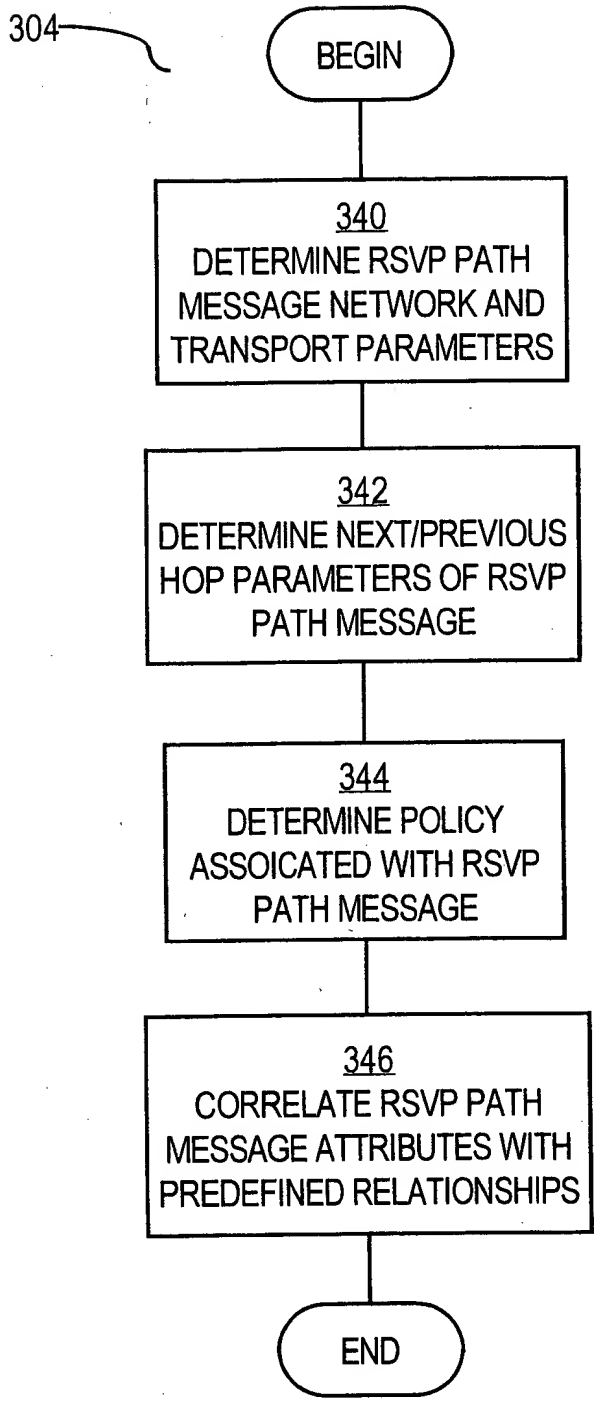
304
BEGIN
340
DETERMINE RSVP PATH MESSAGE NETWORK AND TRANSPORT PARAMETERS
342
DETERMINE NEXT/PREVIOUS HOP PARAMETERS OF RSVP PATH MESSAGE
344
DETERMINE POLICY ASSOICATED WITH RSVP PATH MESSAGE
346
CORRELATE RSVP PATH MESSAGE ATTRIBUTES WITH PREDEFINED RELATIONSHIPS
END

FIG. 6

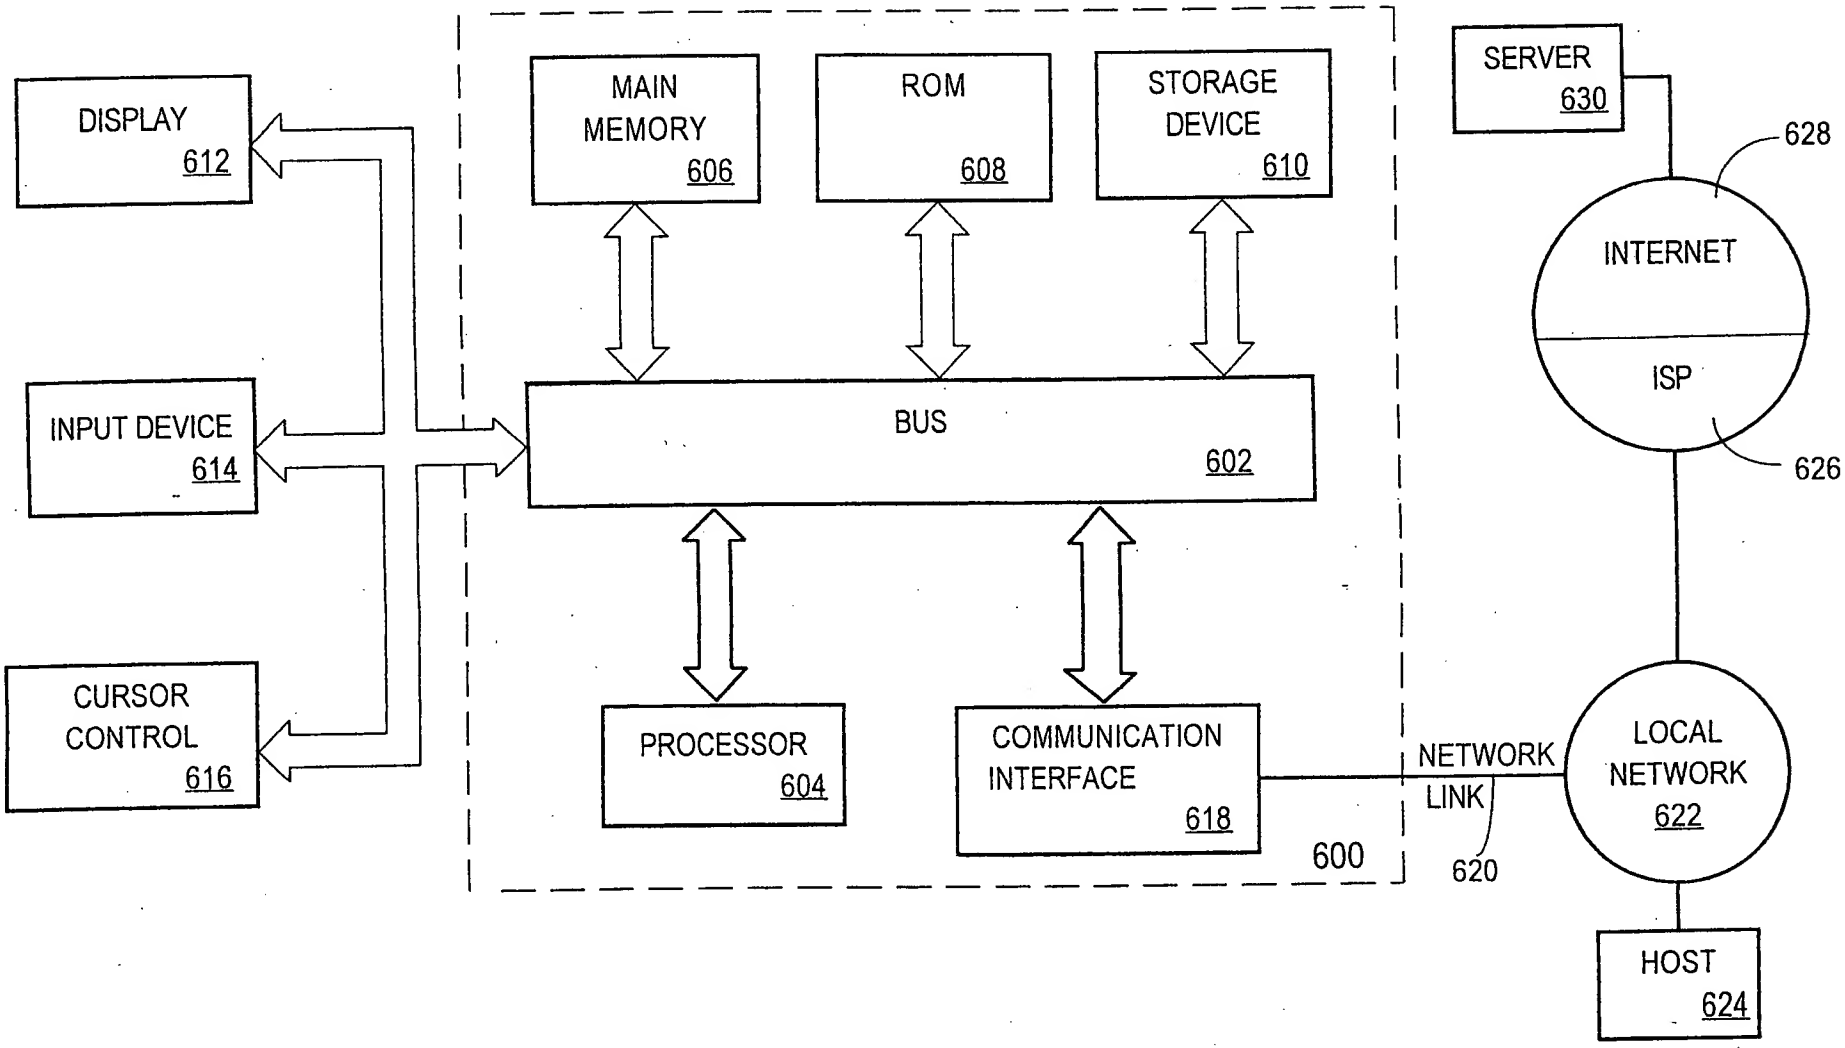
DISPLAY
612
INPUT DEVICE
614
CURSOR
CONTROL
616
MAIN
MEMORY
606
ROM
608
STORAGE
DEVICE
610
BUS
602
PROCESSOR
604
COMMUNICATION
INTERFACE
618
600
NETWORK
LINK
620
SERVER
630
628
INTERNET
ISP
626
LOCAL
NETWORK
622
HOST
624

Currently pending claims of your application

1. A method of establishing a network resources reservation for an anticipated traffic flow along a path in a network between an anticipated source and an anticipated receiver of the traffic flow, wherein the anticipated receiver otherwise cannot facilitate establishing the network resources reservation, the method comprising the steps of:

storing, at the proxy node, policy information defining whether the proxy node should initiate network resources reservations for particular traffic flows;

detecting a first RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a router, acting as a proxy node, located within the path;

determining, at the proxy node and based on the policy information stored at the proxy node and without receiving the policy information from a policy server residing on the network, whether to establish the network resources reservation;

generating, at the proxy node, an RESV message to reserve network resources for the anticipated traffic flow;

communicating the RESV message to the anticipated source of the anticipated traffic flow;

wherein the step of determining, at the proxy node, whether to establish the network resources reservation includes the steps of:

determining one or more network parameter values associated with the anticipated traffic flow;

determining one or more transport parameter values associated with the anticipated traffic flow;

determining next and previous hop parameter values associated with the anticipated traffic flow; and

correlating at least one of the ascertained network parameter, transport parameter, next hop parameter, and previous hop parameter values with information defining a relationship between them and whether a RESV message is desired.

2. A method as recited in claim 1, further comprising the step of determining one or more device and traffic parameter values associated with the anticipated traffic flow, and wherein the step of generating the RESV message comprises the step of generating the RESV message based on at least one of the device and traffic parameter values.

3. (Cancelled).

4. A method as recited in claim 1, further comprising the step of, concurrently with the generating and communicating steps, forwarding a second RSVP Path message to one or more devices that are along the anticipated path and that are between the proxy node and the anticipated receiver, wherein the second RSVP Path message defines a different set of traffic characteristics for the flow initiated by the sender than the first RSVP message.

5. A method as recited in claim 1, wherein determining the network parameter values and ascertaining the transport parameter values includes the steps of determining at least one of the source and receiver IP addresses, source and receiver port numbers, and transport protocol based on values carried in objects in the first RSVP Path message.

6. A method as recited in claim 1, wherein determining the anticipated traffic flow characteristics includes determining at least one of the rate and size of packets associated with the anticipated traffic flow.

7. A method as recited in claim 1, further comprising the steps of extracting one or more additional anticipated traffic flow attributes from the first RSVP Path message.

8. A method as recited in claim 7, wherein the anticipated receiver is an IP phone, and further comprising the step of determining at least one quality of service parameter associated with the anticipated traffic flow.

9. (Canceled)

10. A method as recited in claim 1, wherein the step of detecting an RSVP Path message comprises the step of detecting the first RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a proxy node that is logically positioned adjacent to the path.

11. A computer readable medium comprising one or more sequences of instructions for facilitating an RSVP reservation process, for an anticipated traffic flow anticipated to be received by an anticipated receiver that cannot facilitate an RSVP reservation process for the anticipated traffic flow, wherein when the instructions are executed by one or more processors, the instructions cause the one or more processors to carry out the steps of:

storing, at the proxy node, policy information defining whether the proxy node should initiate network resources reservations for particular traffic flows;

detecting a first RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a router, acting as a proxy node, located within the path;

determining, at the proxy node and based on the policy information stored at the proxy node and without receiving the policy information from a policy server residing on the network, whether to establish the network resources reservation;

generating, at the proxy node, an RESV message to reserve network resources for the anticipated traffic flow;

communicating the RESV message to the anticipated source of the anticipated traffic flow;

wherein the step of determining, at the proxy node, whether to establish the network resources reservation includes the steps of:

determining one or more network parameter values associated with the anticipated traffic flow;

determining one or more transport parameter values associated with the anticipated traffic flow;

determining next and previous hop parameter values associated with the anticipated traffic flow; and

correlating at least one of the ascertained network parameter, transport parameter, next hop parameter, and previous hop parameter values with information defining a relationship between them and whether a RESV message is desired.

12. A computer-readable medium as recited in claim 11, further comprising the step of determining one or more device and traffic parameter values associated with the anticipated traffic flow, and wherein the step of generating the RESV message comprises the step of generating the RESV message based on at least one of the device and traffic parameter values.

13. (Cancelled).

14. A computer-readable medium as recited in claim 11,
further comprising the steps of, concurrently with the generating and communicating steps, forwarding a second RSVP Path message to one or more devices that are along the anticipated path and that are between the proxy node and the anticipated receiver, wherein the second RSVP Path message defines a different set of traffic characteristics for the flow initiated by the sender than the first RSVP message.

15. A computer-readable medium as recited in claim 11, wherein determining the network parameter values and ascertaining the transport parameter values includes the steps of determining at least one of the source and receiver IP addresses, source and receiver port numbers, and transport protocol based on values carried in objects in the first RSVP Path message.

16. A computer-readable medium as recited in claim 11, wherein determining the anticipated traffic flow characteristics includes determining at least one of the rate and size of packets associated with the anticipated traffic flow.

17. A computer-readable medium as recited in claim 11, further comprising the steps of extracting one or more additional anticipated traffic flow attributes from the first RSVP Path message.

18. A computer-readable medium as recited in claim 17, wherein the anticipated receiver is an IP phone, and further comprising the step of determining at least one quality of service parameter associated with the anticipated traffic flow.

19. (Canceled)

20. A computer-readable medium as recited in claim 11, wherein the step of detecting an RSVP Path message comprises the step of detecting the first RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a proxy node that is logically positioned adjacent to the path.

21. A system for establishing a network resources reservation for an anticipated traffic flow along a path in a network between an anticipated source and an anticipated receiver of the traffic flow, wherein the anticipated receiver otherwise cannot facilitate establishing the network resources reservation, the system comprising:

means for storing, at the proxy node, policy information defining whether the proxy node should initiate network resources reservations for particular traffic flows;

means for detecting a first RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a router, acting as a proxy node, located within the path;

means for determining, at the proxy node and based on the policy information stored at the proxy node and without receiving the policy information from a policy server residing on the network, whether to establish the network resources reservation;

means for generating, at the proxy node, an RESV message to reserve network resources for the anticipated traffic flow;

means for communicating the RESV message to the anticipated source of the anticipated traffic flow; and

wherein the means for determining, at the proxy node, whether to establish the network resources reservation includes:

means for determining one or more network parameter values associated with the anticipated traffic flow;

means for determining one or more transport parameter values associated with the anticipated traffic flow;

means for determining next and previous hop parameter values associated with the anticipated traffic flow; and

means for correlating at least one of the ascertained network parameter, transport parameter, next hop parameter, and previous hop parameter values with information defining a relationship between them and whether a RESV message is desired.

22. A network device that can establish a network resources reservation for an anticipated traffic flow along a path in a network between an anticipated source and an anticipated receiver of the traffic flow, wherein the anticipated receiver otherwise cannot facilitate establishing the network resources reservation, the network device comprising:

a network interface;

a processor coupled to the network interface and receiving network messages from the network through the network interface;

a computer-readable medium comprising one or more stored sequences which, when executed by the processor, cause the processor to carry out the steps of:

storing, at the proxy node, policy information defining whether the proxy node should initiate network resources reservations for particular traffic flows;

detecting a first RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a router, acting as a proxy node, located within the path;

determining, at the proxy node and based on the policy information stored at the proxy node and without receiving the policy information from a policy

server residing on the network, whether to establish the network resources reservation;

generating, at the proxy node, an RESV message to reserve network resources for the anticipated traffic flow;

communicating the RESV message to the anticipated source of the anticipated traffic flow; and

wherein the step of determining, at the proxy node, whether to establish the network resources reservation comprises the steps of:

determining one or more network parameter values associated with the anticipated traffic flow;

determining one or more transport parameter values associated with the anticipated traffic flow;

determining next and previous hop parameter values associated with the anticipated traffic flow; and

correlating at least one of the ascertained network parameter, transport parameter, next hop parameter, and previous hop parameter values with information defining a relationship between them and whether a RESV message is desired.

23. A system as recited in claim 21, further comprising means for determining one or more device and traffic parameter values associated with the anticipated traffic flow, and wherein the means for generating the RESV message comprises means for generating the RESV message based on at least one of the device and traffic parameter values.

24. (Cancelled).

25. A system as recited in claim 21,

further comprising means for forwarding, concurrently with operation of the means for generating and the means for communicating, a second RSVP Path message to one or more devices that are along the anticipated path and that are between the proxy node and the anticipated receiver, wherein the second RSVP Path message

defines a different set of traffic characteristics for the flow initiated by the sender than the first RSVP message.

26. A system as recited in claim 24, wherein the means for determining the network parameter values and ascertaining the transport parameter values includes means for determining at least one of the source and receiver IP addresses, source and receiver port numbers, and transport protocol based on values carried in objects in the first RSVP Path message.

27. A system as recited in claim 24, wherein the means for determining the anticipated traffic flow characteristics includes means for determining at least one of the rate and size of packets associated with the anticipated traffic flow.

28. A system as recited in claim 24, further comprising means for extracting one or more additional anticipated traffic flow attributes from the first RSVP Path message.

29. A system as recited in claim 27, wherein the anticipated receiver is an IP phone, and further comprising means for determining at least one quality of service parameter associated with the anticipated traffic flow.

30. A system as recited in claim 21, wherein the means for detecting an RSVP Path message comprises means for detecting a first RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a proxy node that is logically positioned adjacent to the path.

31. A network device as recited in claim 22, wherein the one or more stored sequences, when executed by the processor, cause the processor to further carry out the step of determining one or more device and traffic parameter values associated with the anticipated traffic flow, and wherein the step of generating the RESV message comprises the step of

generating the RESV message based on at least one of the device and traffic parameter values.

32. (Cancelled).

33. A network device as recited in claim 22,
further comprising instructions for performing the step of, concurrently with the generating and communicating steps, forwarding a second RSVP Path message to one or more devices that are along the anticipated path and that are between the proxy node and the anticipated receiver, wherein the second RSVP Path message defines a different set of traffic characteristics for the flow initiated by the sender than the first RSVP message.

34. A network device as recited in claim 22, wherein determining the network parameter values and ascertaining the transport parameter values includes the steps of determining at least one of the source and receiver IP addresses, source and receiver port numbers, and transport protocol based on values carried in objects in the first RSVP Path message.

35. A network device as recited in claim 22, wherein determining the anticipated traffic flow characteristics includes determining at least one of the rate and size of packets associated with the anticipated traffic flow.

36. A network device as recited in claim 22, wherein the one or more stored sequences, when executed by the processor, cause the processor to further carry out the step of extracting one or more additional anticipated traffic flow attributes from the RSVP Path message.

37. A network device as recited in claim 36, wherein the anticipated receiver is an IP phone, and wherein the one or more stored sequences, when executed by the processor, cause the processor to further carry out the step of determining at least one quality of service parameter associated with the anticipated traffic flow.

38. A network device as recited in claim 22, wherein the step of detecting an RSVP Path message comprises the step of detecting the first RSVP Path message associated with the anticipated receiver of the anticipated traffic flow at a proxy node that is logically positioned adjacent to the path.

Network Working Group
Internet Draft
draft-sgai-rsvp-proxy-00.txt
Expiration Date: April 2000

Silvano Gai
Dinesh Dutt
Nitsan Elfassy
Cisco Systems
Yoram Bernet
Microsoft

October 1999

RSVP Receiver Proxy

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt. The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

Distribution of this memo is unlimited.

Copyright Notice

Abstract

RSVP has been extended in several directions [Policy], [Identity], [DCLASS], [AggrRSVP], [DiffModel], [COPS-RSVP], These extensions have broadened the applicability of RSVP characterizing it as a signaling protocol usable outside the IntServ model.

With the addition of the "Null Service Type" [NullServ], RSVP is being adopted also by mission critical applications that require some form of prioritized service, but cannot readily specify their resource requirements. These applications do not need to set-up a reservation end-to-end, but only to signal to the network their policy information [Policy], [Identity] and obtain in response an applicable DSCP [DCLASS].

RSVP Receiver Proxy is an extension to the RSVP message processing (not to the protocol itself), mainly designed to operate in conjunction with the Null Service Type and with an extension of the COPS for RSVP protocol [COPS-RSVP-EXT].

Table of contents

1. Introduction

The IETF has come up with two architectures to support QoS in IP networks. IntServ (Integrated Services [RFC1633], [RFC2210]) is an architecture that provides the ability for applications to choose among multiple, controlled levels of delivery service for their data packets. It relies upon explicit signaling by applications to the network for the desired QoS. These applications typically know their traffic characteristics and have possibly strict latency requirements. Such applications require so called "tight QoS" or "quantitative QoS". RSVP is the protocol which can be used by applications to signal their QoS requirements to the network. Applications have to be modified to take advantage of the Integrated Services. The receivers control the QoS given to the data stream.

DiffServ (Differentiated Services, [RFC2474], [RFC2475]) is another IETF architecture for implementing scalable service differentiation

in the Internet. There is no explicit signaling protocol used in DiffServ. The network is logically divided into edge devices and core devices. The edge devices attempt to recognize data flows and assign QoS based on this. They also assign a DSCP (DiffServ Code Point) in the DS byte of the packets (the byte that used to be called the TOS byte). Core devices use the DSCP to assign a QoS to the microflows. Applications typically do not have to be modified to take advantage of Differentiated Services. Receivers do not control the QoS given to the data stream.

The recognition of data flows and the assignment of an appropriate DSCP is a tricky task and often requires stateful inspection of flows and symmetrical routing paths. Moreover, application recognition is limited to the information present in the packet traversing the network and in most current network devices is further limited to what is in the IP/TCP/UDP headers. Application vendors desire to be able to assign QoS to their packets based on both information that may not be carried in the packet and information other than the IP/TCP/UDP header fields. For example, a SAP print transaction may require a different treatment than a SAP database update. Similarly, if the user of the application is the CTO of the company, the priority assigned to such packets maybe different from that assigned to packets of the application being used by some other person in the company.

For this reason RSVP has been proposed also for mission critical applications (e.g. ERP) that require some form of prioritized service, but cannot readily specify their resource requirements. The ISSLL WG is discussing the specification of the Null Service Type as a way to use RSVP with a broader range of applications [NullServ].

Some of these applications have the requirement for the end-to-end message processing of RSVP. Others simply need to signal to the network their identity [Identity] and some additional policy information [Policy] related to the flows and obtaining from the network some decisions, e.g. the DSCP to be used [DCLASS].

RSVP Receiver Proxy is a proposal that mainly addresses this second type of applications, i.e., applications that simply want to use RSVP as a signaling protocol toward the network. For them, the end-to-end nature of RSVP is not interesting and often is perceived as a disadvantage, since it is characterized by a higher latency.

The RSVP Receiver Proxy:

- is an alternate way to process RSVP messages and policy information in the switch/routers;
- it does not require any change to the RSVP protocol;
- it does require an extension to the COPS for RSVP protocol [COPS-RSVP-EXT].

In general, "RSVP Proxy" should be symmetric, i.e., it may be useful to have RSVP Sender Proxy as well as RSVP Receiver Proxy. This document does not define RSVP Sender Proxy at this stage. If the document is accepted by the IETF community, the RSVP Sender Proxy can

be added in the next version.

This document defines RSVP Receiver Proxy in association with the Null Service Type, but nothing prevents using this feature also in association with other service types, e.g. the Controlled Load service.

The following section uses an example in which the Receiver Proxy functionality is placed in the first hop switch/router. This is a possibility, but it is not a requirement. While designing a network the following trade-off should be considered:

- Proxying closer to the server reduces turn around time.
- Proxying further from the server enables additional downstream network elements to benefit from the information carried in the signaling messages, and to participate in the response.
- Proxying anywhere in the network enables the deployment of such applications in which only the server is required to signal, but

the client may remain unchanged.

The COPS-RSVP Extension [COPS-RSVP-EXT] should enable the network administrator to decide how to make the tradeoffs described above.

2. An overview of RSVP Receiver Proxy

With RSVP Receiver Proxy a switch/router acts as a proxy for the receiver, e.g. when it receives an RSVP Path message, it generates an RSVP Resv message on behalf of the receiver.

The generation of the Resv message is done under policy control, the switch/router may be programmed either to classify the packets marking them with an appropriate DSCP or to use the DCLASS object [DCLASS] to communicate the classification decision to the host.

The adoption of RSVP Receiver Proxy do not change the basic model of RSVP, i.e.:

- the handling of data flows is unidirectional. If the application data is strictly unidirectional it is sufficient to use RSVP only in one direction. In the case of bidirectional data, running RSVP only in one direction provides a certain performance benefit, but to get the maximum performance benefit it is necessary to use RSVP in both directions.
- The application on the host assumes the host model of RSVP, including the extensions proposed in [DiffModel], [Policy], [Identity], [NullServ].
- The message format and the message types are the same of RSVP, including the DCLASS object previously proposed in [DCLASS] and the Null Service Type [NullServ].
- The switch/router acts as a COPS client [COPS] in communicating with the policy server, i.e. it uses RSVP client for COPS [COPS-

RSVP]. Certain extensions to COPS for RSVP are needed [COPS-RSVP-EXT], see Section 4.

- o The classification of traffic cannot be more granular than microflow (the so called five-tuple) or in the case of IPSEC the four-tuple that includes the Parameter Index, or SPI, in place of the UDP/TCP-like ports [RFC2207].

- o There is no special support for subflows (a set of packets inside a microflow). Of course, an application may send different Path

messages for the same flow at different times, thus providing a support for subflows not overlapping in time.

3. Detailed description of the message processing

This sections details some of the message processing of a switch/router acting as RSVP Receiver Proxy. The description is mainly focused on the two fundamental messages in RSVP, i.e. the Path Message and the Resv message. Other messages are discussed in Section 4.4.

Figure 1 depicts a simple network topology (two hosts H1 & H2 and intermediate routers, R1-R5) that will be used in the explanation.

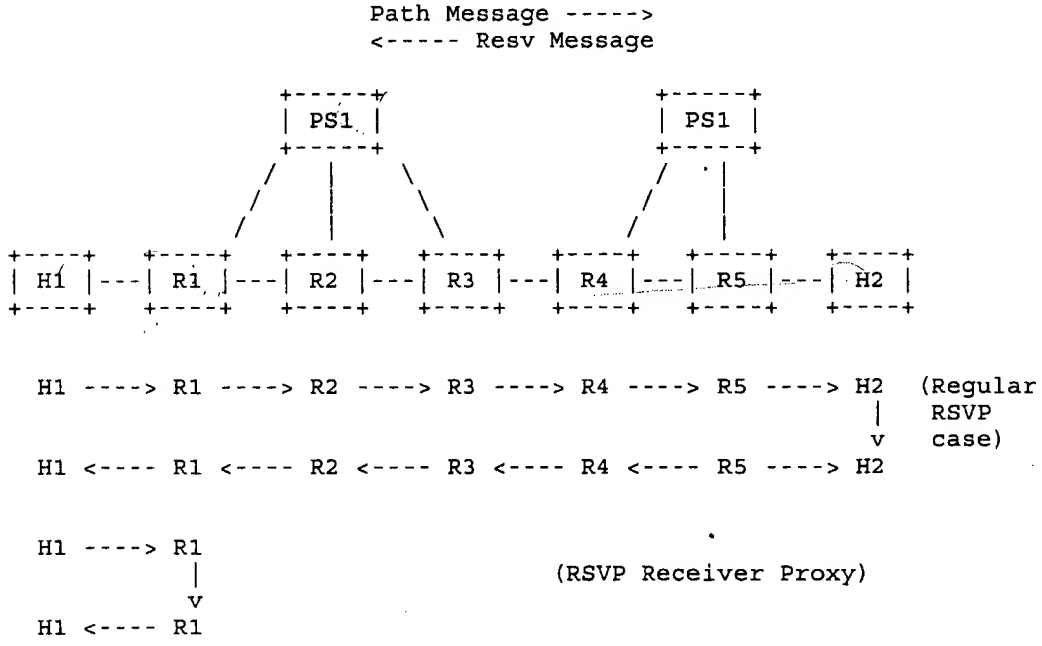

Hx: Host x
Ry: Router y
PSz: Policy Server z

Figure 1: Possible Message Forwarding Behaviors in RSVP

Immediately below the network, the normal RSVP message processing is reported. The Path message goes hop-by-hop from H1 to H2. The Resv message uses the reverse path of the Path message and goes from H2 to H1. The interaction between the network devices and the policy servers is the one specified by COPS for RSVP ([COPS], [COPS-RSVP]).

With RSVP Receiver Proxy the propagation of the RSVP Path message is terminated in the router acting as a proxy. Any router in the network may act as RSVP Receiver Proxy, but it is a good design guideline to place the proxy functionality as close as possible to the sender. In our case R1 acts as a proxy for H2 under the control of a policy server.

For example, an application on H1 uses RSVP to signal parameters upon which to base the decision to assign the QoS for a microflow. The example assumes that the information needs to be used only by the edge network device and it is not required to propagate this further down the network

A possible sequence of steps consists of:

- The application on H1 indicates to the RSVP subsystem that it is a sender and specifies its traffic characteristics. It may specify additional parameters.

- This causes the RSVP subsystem on H1 to start transmitting RSVP Path messages in accordance with normal RSVP/SBM rules.

- The first hop switch/router (R1) receives this message and it communicates with the policy server for a decision on how to treat the Path message. It copies all the relevant information contained in the Path message to the policy server.

- The policy server communicates a decision to R1 to not forward the Path message, but instead to originate and send a Resv message to H1. H1 data traffic gets assigned the right DSCP by the switch/router as per the policy communicated by the policy server. The Resv message may also specify to the host the DSCP and shaping information to be associated with the microflow using the DCLASS object [DCLASS].

- On receiving the Resv message, H1 may start marking correctly the data traffic accordingly to the DSCP received in the Resv message.

4. The role of the policy server

To implement both RSVP and RSVP Receiver Proxy the policy server needs to specify a set of decisions [COPS-RSVP-EXT] which is extended compared to COPS-RSVP [COPS-RSVP]. If the decision is to accept the Path message, the decision message must specify how the network device behaves with respect to each of the following:

- o Forwarding of the Path message;
- o Originating a RSVP Resv message;
- o Processing and possibly Forwarding a RSVP Resv message.

The decision may also possibly include the QoS specification to be associated with the flow identified in the Path message. This specification consists of a DSCP and possibly a TSPEC (as specified by RSVP [RFC2210]) for policing the traffic.

4.1 Generation of the Resv message by the Receiver Proxy

It maybe required that the network device originate a Resv message. This is a proxy Resv message in the sense that it is being generated by the network device and not by the actual receiver(s) identified in the RSVP Path message. The format of a Resv message is as follows (see [RFC2205] for details):

```
<Resv Message> ::=      <Common Header> [ <INTEGRITY> ]
   <SESSION> <RSVP HOP> <TIME_VALUES><DCLASS>
   [ <RESV_CONFIRM> ] [ <SCOPE> ] [ <POLICY_DATA>... ]
   <STYLE> <flow descriptor list>
```

- o The network device puts its IP address and L2 address in the source IP and source mac-address fields. Since Resv messages follow Path messages, this would constitute a valid Resv message.
- o The SESSION object can be copied from the Path message.
- o The RSVP HOP object can be filled in with the IP address of the switch/router generating this Resv message.

- o The TIME_VALUES object contains the refresh period. See below.
- o The STYLE object is set to Wildcard Filter (WF) style indicating that the reservation is to be shared and that the sender is wildcarded. Associated with a WF style is a FLOWSPEC object which is encoded as specified in [RFC2210] or [NullServ].
- o The SCOPE and RESV_CONFIRM objects need not be included in the Resv message.

- The POLICY_DATA objects will be as returned by the policy server.

- The Resv message may also contain the new DCLASS object is contained in the COPS decision message. The DCLASS object specifies the DSCP to be associated with the microflow for which the Path message was received.

- The Resv messages need to be originated and sent for each of the periodically-received Path messages.

4.2 Communication With the Policy Server

When a network device establishes the connection with the policy server, it sends a COPS Client-Open message for the RSVP client. It should indicate in this message whether the network device is capable of supporting only the base RSVP message processing or also the Receiver Proxy message processing. It can do this with in a capability list (that can accommodate also future extensions). To deal with existing clients, if the policy server does not receive a capability list, it should assume that it is communicating with a legacy RSVP client. The capability list can be included as part of the ClientSI object passed in the Client-Open message [COPS-RSVP-EXT].

On receiving a RSVP Path message, the network device sends a COPS REQ message to the policy server. This message will be the standard REQ message sent on receiving a RSVP Path message.

The DEC message returned by the policy server for this REQ message must contain the information needed to take the decisions listed in Section 4.

The DEC message SHOULD also contain a list of DSCP [DCLASS].

The DEC message may also contain bandwidth information to be associated with the microflow: communicating Shaping/limiting

parameters to the network is a powerful Policy Management tool for the PDP/LPDP both for Qualitative and Quantitative services. This topic needs further study.

The network device must also be able to determine if a Path message is a refresh or a new one. It must communicate with the policy server only for new Path messages or for updated ones.

In the absence of a policy server or if the connection to the policy server is not up, the operation of RSVP Receiver Proxy depends on policy configuration local to the network device. For example, the network device may have a local configuration that specifies:

- do not accept new flows;

- honor existing flows until they time-out.

4.3 Enhancements To Existing Infrastructure

- o COPS for RSVP will have to be enhanced to support the new format for RSVP REQ and DEC message as stated in [COPS-RSVP-EXT].

- o When SBM is in use, it is possible that a device which does not support RSVP Receiver Proxy becomes the DSBM on the first-hop segment. This can be prevented by the network administrator by configuring the appropriate priority on the device with RSVP Receiver Proxy support.

4.4 Processing of other RSVP messages

This section details the processing of the protocol messages in RSVP other than Path and Resv. Only the differences in the processing from classical RSVP is specified.

- o PathTear message is honored and is forwarded or not similar to a Path message. The policy server is not contacted on receiving a PathTear message. This is consistent with the existing behavior of COPS for RSVP[RSVP-COPS].

- o PathErr messages are treated as in normal RSVP.

5. RSVP With Null Service Type

RSVP protocol can be represented as consisting of two parts: a message processing part and a resource allocation & resource enforcement part. The following are the minimal requirements for a network device to support RSVP Null Service Type:

- o The network device MUST implement the message processing part of the RSVP protocol. This includes the ability to receive and interpret a raw IP packet or UDP-based RSVP packet.

- o If the network device is a L2 device, it SHOULD implement SBM.

- o The network device SHOULD know how to talk to a policy server using COPS. Specifically, the network device SHOULD be able to talk to COPS as a RSVP client using the extensions defined in [COPS-RSVP-EXT].

- o The node SHOULD keep the RSVP state so that the following Path refresh won't cause a repetitive Path handling.

- o The network device SHOULD be able to generate a Resv message periodically in a coherent way with the RSVP soft state maintenance.

- o In the absence of a connection to the policy server, this network device depends on policy configuration local to the network device (see Section 4.2).

6. Security Considerations

RSVP messages contain an INTEGRITY object which authenticates the originating node and is also used to verify the contents of the message. Moreover the RSVP message SHOULD contain an IDENTITY object that SHOULD be authenticated. If the policy server does not implement any security mechanisms, it SHOULD use a clear text version of the user identity.

7. Intellectual Property Considerations

The IETF is being notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

8. References

[COPS] Boyle, J., Cohen, R., Durham, D., Herzog, S., Raja, R., Sastry, A., "The COPS (Common Open Policy Service) Protocol", IETF <draft-ietf-rap-cops-07.txt>, August 1999.

[RFC1633] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet Architecture: an Overview," June 1994.

[RFC2205] Braden, R., Zhang, L., Berson, S., Herzog, S., and Jamin, S., "Resource Reservation Protocol (RSVP) Version 1 Functional Specification", IETF RFC 2205, Proposed Standard, September 1997.

[RFC2210] J. Wroclawski, "The Use of RSVP with IETF Integrated Services," September 1997.

[RFC2474] K. Nichols, S. Blake, F. Baker, D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," December 1998.

[RFC2475] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, "An Architecture for Differentiated Service," RFC 2475, December 1998.

[COPS-RSVP] Jim Boyle, Ron Cohen, David Durham, Shai Herzog, Raju Rajan, Arun Sastry, "COPS usage for RSVP," <draft-ietf-rap-cops-rsvp-05.txt>, June 14, 1999

[COPS-RSVP-EXT] Nitsan Elfassy, Dinesh Dutt, "COPS Extensions for RSVP Receiver Proxy Support"," <draft-nitsan-cops-rsvp-proxy-00.txt>, October 1999.

[Policy] Shai Herzog, "RSVP Extensions for Policy Control," Internet Draft., < draft-ietf-rap-rsvp-ext-06.txt>, April 1999.

[DiffModel] Y. Bernet, A. Smith, S. Blake, "A Conceptual Model for

Diffserv Routers," Internet Draft, <draft-ietf-diffserv-model-00.txt>, June 1999.

[Identity] Satyendra Yadav, Raj Yavatkar, Ramesh Pabbati, Peter Ford, Tim Moore, Shai Herzog, "Identity Representation for RSVP," Internet-Draft <draft-ietf-rap-rsvp-identity-05.txt>, September 1999.

[AggrRSVP] Fred Baker, Carol Iturralde, Francois Le Faucheur, Bruce Davie, "Aggregation of RSVP for IP4 and IP6 Reservations," <draft-ietf-issll-rsvp-aggr-00.txt>, September 1999

[DCLASS] Bernet, Y., "Usage and Format of the DCLASS Object With RSVP Signaling," <draft-ietf-issll-dclass-00.txt >, August 1999.

[NullServ] Yoram Bernet, Andrew Smith, B. Davie, "Specification of the Null Service Type," <draft-ietf-issll-nullservice-00.txt>, September 1999

[RSVPDIFF] Bernet, R. Yavatkar, P. Ford, F. Baker, L. Zhang, M. Speer, B. Braden, B. Davie, J. Wroclawski, E. Felstaine, "Integrated Services Operation Over Diffserv Networks," <draft-ietf-issll-diffserv-rsvp-03.txt>, September 1999

9. Author Information

Silvano Gai
Cisco Systems, Inc.
170 Tasman Dr.
San Jose, CA 95134-1706
Phone: (408) 527-2690
email: sgai@cisco.com

Dinesh Dutt
Cisco Systems, Inc.
170 Tasman Dr.
San Jose, CA 95134-1706
Phone: (408) 527-0955
email: ddutt@cisco.com

Nitsan Elfassy
Cisco Systems, Inc.
Cisco Systems, Inc.
170 Tasman Dr.
San Jose, CA 95134-1706
Phone: +972 9 970 0066
email: nitsan@cisco.com

Bernet, Yoram
Microsoft
One Microsoft Way,
Redmond, WA 98052
Phone: (425) 936-9568
Email: yoramb@microsoft.com

10. Full Copyright Statement

Hickman Palermo Truong & Becker, LLP
2055 Gateway Place, Suite 550
San Jose, CA 95110-1089

RETURN TO SENDER
NUMBER
NOT KNOWN
ROUTE NUMBER 1629

PLACE STICKER AT TOP OF ENVELOPE TO THE RIGHT OF RETURN ADDRESS FOLD AT DOTTED LINE
CERTIFIED MAIL

7000 1670 0006 8330 2896

10/29 SCANNED
RET. NOT KNOWN

Angelica

UNITED STATES POSTAGE
PB946742 1
185
1594 $06.72 OCT 28 05
2653 MAILED FROM ZIP CODE 95125

SENDER: *COMPLETE THIS SECTION*

- Complete items 1, 2, and 3. Also complete item 4 if Restricted Delivery is desired.
- Print your name and address on the reverse so that we can return the card to you.
- Attach this card to the back of the mailpiece, or on the front if space permits.

1. Article Addressed to:

Itzhak Parnafes
20050 Rodrigues Avenue, Apt. B
Cupertino, CA 95014

COMPLETE THIS SECTION ON DELIVERY

A. Received by *(Please Print Clearly)* | B. Date of Delivery

C. Signature
X ☐ Agent ☐ Addressee

D. Is delivery address different from item 1? ☐ Yes
If YES, enter delivery address below: ☐ No

3. Service Type
☒ Certified Mail ☐ Express Mail
☐ Registered ☐ Return Receipt for Merchandise
☐ Insured Mail ☐ C.O.D.

4. Restricted Delivery? *(Extra Fee)* ☐ Yes

2. Article Number *(Copy from service label)*
7000 1670 0006 8330 2896

PS Form 3811, July 1999 Domestic Return Receipt 102595-99-M-1789

RETURN RECEIPT REQUESTED

BEST AVAILABLE COPY

HICKMAN PALERMO TRUONG & BECKER, LLP
INTELLECTUAL PROPERTY LAW
2055 GATEWAY PLACE, SUITE 550
SAN JOSE, CALIFORNIA 95110-1089

PLACE STICKER AT TOP OF ENVELOPE TO THE RIGHT OF RETURN ADDRESS. FOLD AT DOTTED LINE

CERTIFIED MAIL

7000 1670 0006 8330 3381

UNITED STATES POSTAL SERVICE

0000 94303

U.S. POSTAGE PAID SANJOSE,CA 95110 MAR 29, 06 AMOUNT **$0.00** 00000576-20

RETURN TO SENDER

- ☐ INSUFFICIENT ADDRESS
- ☐ NO SUCH NUMBER
- ☐ UNCLAIMED ☐ REFUSED
- ☑ ATTEMPTED NOT KNOWN
- ☐ NO SUCH STREET
- ☐ VACANT
- ☐ NO RECEPTACLE
- ☐ NOT DELIVERABLE AS ADDRESSED - UNABLE TO FORWARD

ROUTE NO. ____ DATE ____
ABBR/INITIALS ____

UNITED STATES POSTAGE PB93502604
★ ★ ★ 179 4791 $03.27 0 MAR 29 06 3512 MAILED FROM ZIP CODE 95110

UNITED STATES POSTAGE PB93502604
★ ★ ★ 186 4741 $04.25 0 MAR 29 06 3513 MAILED FROM ZIP CODE 95110

SENDER: *COMPLETE THIS SECTION*

- Complete items 1, 2, and 3. Also complete item 4 if Restricted Delivery is desired.
- Print your name and address on the reverse so that we can return the card to you.
- Attach this card to the back of the mailpiece, or on the front if space permits.

1. Article Addressed to:

Itzhak Parnafes
925 Piers Court
Palo Alto, CA 94303

COMPLETE THIS SECTION ON DELIVERY

A. Received by *(Please Print Clearly)* B. Date of Delivery

C. Signature

X ☐ Agent ☐ Addressee

D. Is delivery address different from item 1? ☐ Yes
If YES, enter delivery address below: ☐ No

3. Service Type
☑ Certified Mail ☐ Express Mail
☐ Registered ☐ Return Receipt for Merchandise
☐ Insured Mail ☐ C.O.D.

4. Restricted Delivery? *(Extra Fee)* ☐ Yes

2. Article Number *(Copy from service label)*
7000 1670 0006 8330 3381

PS Form 3811, July 1999 Domestic Return Receipt 102595-99-M-1789

RETURNED
NO FORWARD
ON FILE

BEST AVAILABLE COPY